An exploration of spatio-temporal statistical models for landslide hazard assessment and prediction

Erin Bryce

Submitted in fulfilment of the requirements
for the Degree of Doctor of Philosophy

School of Mathematics and Statistics
College of Science and Engineering
University of Glasgow

August 2025

Declaration of Authorship

I, Erin Bryce, do hereby certify that this thesis, submitted for the degree of PhD, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for any degree.

I was admitted as a research student at the University of Glasgow in October 2021 and this thesis was developed wholly while under candidature for a PhD at the University of Glasgow.

I received funding support by the Additional Funding Programme for Mathematical Sciences, delivered by EPSRC (EP/V521917/1) and the Heilbronn Institute for Mathematical Research.

The work in Chapter 2 was carried out in collaboration with my supervisors, Dr. Daniela Castro-Camilo at the University of Glasgow and Dr. Luigi Lombardo at ITC, University of Twente (The Netherlands), alongside Dr. Hakan Tanyas and Prof. Cees van Westen - both also at ITC. The results of which were presented as a talk at the International Association for Mathematical Geosciences in 2022, in Nancy, France and published in the journal *Stochastic Environmental Research and Risk Assessment*.

The work in Chapter 3 was in collaboration with Dr. Alessandro Novellino, Dr. Roxana Ciurean and Claire Dashwood at the British Geological Survey, my supervisors; Dr. Daniela Castro-Camilo and Dr. Luigi Lombardo, and Dr. Hakan Tanyas. This work was presented as a talk at the Research Students' Conference in Probability and Statistics in 2023, in Sheffield, UK, as a poster at ENRA Science, Evidence and Policy Conference in 2024, in Edinburgh, UK, and published in the journal *Landslides*.

The work in Chapter 4 was carried out in collaboration with Prof. Sara Martino at (Norwegian University of Science and Technology), my supervisor Prof. Janine Illian (University of Glasgow), Virginia Morera-Pujol (University College Dublin) and Leonardo Capitani (Swiss Federal Institute for Forest, Snow and Landscape Research, WSL). A manuscript based on this work is currently in preparation. Building on our joint research, I developed an additional case study aligned with my broader research aims. The results are presented in Section 4.5 and have been shared through a talk at the Royal Statistical Society Conference 2024 in Brighton, UK, and poster presentations at the METMA 2024 conference in Lancaster, UK, and the second edition of the workshop INLA: past, present, and future in 2025, held in Glasgow, UK.

The work presented in Chapter 5 was conducted in collaboration with my supervisors, Dr. Daniela Castro-Camilo, Prof. Janine Illian, and Dr. Luigi Lombardo. A manuscript based on this work is currently in preparation.

Abstract

This thesis develops and applies statistical modelling techniques for complex spatial and spatio-temporal geophysical datasets, with a particular focus on the estimation and assessment of landslide hazard and surface deformation. The bulk of this thesis's methodological framework is grounded in Bayesian inference, utilising the integrated nested Laplace approximation (INLA) for efficient computation. The models employed are of the latent Gaussian class, wherein observations are conditioned on an unobserved latent field that captures residual spatial or spatio-temporal variation. This latent field is represented via a Matérn Gaussian Random Field, approximated through the Stochastic Partial Differential Equation (SPDE) approach. Domain-specific covariates - geographical, geological and meteorological - are incorporated within the hierarchical structure. In doing so, the thesis explores landslide hazard in terms of *where* landslides occur, *when* they occur, and *how large* they are.

The statistical approaches developed span a range of modelling strategies, including susceptibility models (presence/absence), Poisson and log-Gaussian Cox processes (LGCPs), functional generalised additive models (FGAMs), and a custom space-time SPDE smoother implemented within the Mixed GAM Computation Vehicle with Automatic Smoothness Estimation package (`mgcv`), which is a flexible framework for modelling non-linear relationships within GAMs. This enables the integration of high-resolution environmental covariates and a functional precipitation predictor, with various continuous distributions used to model landslide size.

Chapter 2 introduces a unified landslide hazard framework through a Hurdle model, jointly modelling occurrence (via a Bernoulli process) and size (via a log-Gaussian model for planimetric extent). This enables the creation of hazard maps that provide probabilistic estimates of large-event exceedances, along with their associated uncertainty.

Chapter 3 presents an updated landslide susceptibility model for Scotland, developed for the British Geological Survey (BGS). It includes a proposed LGCP extension and provides the first data-driven landslide susceptibility framework for the BGS, benchmarked against their previous heuristic model, GeoSure.

Chapter 4 addresses a key methodological challenge: the influence of mesh resolution and integration scheme in SPDE-based point process models with fine-scale covariates. Motivated by issues encountered in the BGS application, a series of simulation studies explores the effects of mesh specification, culminating in a case study where a marked LGCP is fitted to a Japanese landslide inventory using landslide size as the mark.

Chapter 5 explores the temporal dimension of landslide hazard through a spatio-temporal model of surface deformation in a region of China over a two-month period. This chapter introduces a functional precipitation predictor and transitions from a Bayesian to a frequentist framework, motivated by limitations in earlier precipitation representations. It implements, for the first time, a space-time SPDE Matérn smoother within mgcv, enabling flexible modelling of deformation using high-resolution covariates and functional data analysis techniques.

Acknowledgments

Firstly, I would like to thank my three supervisors: Daniela Castro-Camilo, Janine Illian and Luigi Lombardo. When I began this journey, fresh out of my Bachelor's degree they showed me nothing but unwavering support as we worked together on this research. I am particularly thankful to Daniela, who saw something in me when I was assigned to her dissertation project in 2020 and offered me the opportunity to continue the research we started there. I will always be grateful for this opportunity, and the many more you provided me with over the years I was your student - I could not have asked for a better supervisor. From each of my supervisors, I have learned so much in regard to statistical research, writing (thank you Luigi!), and the application chosen for this work. I could never have completed this work without them.

This work was funded by the Heilbronn Institute for Mathematical Research (HIMR), which made this project possible. For specific chapters of my thesis, I would like to thank, in addition to my supervisors and the HIMR, the following people:

Chapter 3:

Alessandro Novellino, Roxana Ciurean and Claire Dashwood for making me feel so welcome at the British Geological Survey (BGS) offices in Keyworth, Nottingham, during my visits. The geological knowledge imparted whilst I was here was fundamental in my understanding of what the desired model output would actually be used for.

Chapter 4:

Sara Martino, for your patience and support in addressing the numerous questions I posed throughout the course of this work. From our discussion at RSS after my talk until the day my models finally ran (!), I cannot thank you enough. From these discussions and the Spatial Point Processes reading group meetings, set up by Megan Laxton and Janine Illian, the 'mess with the mesh' group was born.

Finally, I want to thank my parents, Colin and Chantel, and my brother Aidan. I never felt alone or stressed with your constant love and support, and I am forever blessed to be part of our family. Also, thanks to Liga, Kara, Linda and Olly for weekend adventures and wonderful trips that have given me so much joy.

Contents

List of Figures

List of Tables

List of software

I was able to complete this research exclusively using freely available software, most of it open-source. This section lists the software used to produce the results in this thesis. Although all these tools themselves depend on other software packages, for simplicity, I mention only the software that I used directly, not their dependencies. I also exclude software I used during exploratory work that did not ultimately contribute to the outputs in this thesis. I used Crib, the Geospatial Computing Platform of the ITC, University of Twente, NL, and Windows 10 Pro version on my university laptop. All of the programming was conducted in R. To interface with R, I utilised RStudio. The following R packages were used in the process of producing this thesis:

- `dplyr` (Wickham *et al.*, 2023)
- `fields` (Douglas Nychka *et al.*, 2021)
- `ggplot2` (Wickham, 2016)
- `gridExtra` (Auguie, 2022)
- `INLA` (Rue *et al.*, 2009)
- `inlabru` (Bachl *et al.*, 2019)
- `INLAspacetime` (Krainski *et al.*, 2025)
- `knitr` (Xie, 2024)
- `lattice` (Sarkar, 2008)
- `latticeExtra` (Sarkar and Andrews, 2022)
- `maptools` (Bivand and Lewin-Koh, 2023)
- `matrixStats` (Bengtsson, 2024)
- `mgcv` (Wood, 2017)
- `pacthwork` (Pedersen, 2024)
- `raster` (Hijmans, 2024)
- `refund` (Goldsmith *et al.*, 2024)

- `rgdal` (Bivand *et al.*, 2023)
- `rgeos` (Bivand and Rundel, 2023)
- `sf` (Pebesma *et al.*, 2018)
- `sp` (Pebesma and Bivand, 2005)
- `spdep` (Pebesma and Bivand, 2023)
- `tidyr` (Wickham *et al.*, 2024)
- `viridis` (Garnier *et al.*, 2024)

Chapter 1

Introduction

This thesis explores statistical methodologies for spatial and spatio-temporal modelling of landslide hazard. Landslide hazard is originally conceived as a function of the spatial propensity of landslide events (susceptibility) and the temporal frequency at which they occur (Guzzetti *et al.*, 1999). This definition was later modified in order to account for the intensity associated to a population of landslides (Crozier, 2005). To put it simply, landslide hazard refers to the *where*, the *when* and the *how large*. Each of these aspects can be defined and measured in various ways, and in this work, we explore statistical-based modelling approaches to represent the three aspects of landslide hazard, independently and jointly.

Landslide hazard assessment is crucial in providing a comprehensive evaluation for mitigation measures and evacuation plans. Perhaps one of the most notable landslide hazard risk management schemes in place is that of the Slope Safety System in Hong Kong, China. With an annual average of 2300mm of rainfall a year, and extreme rainfall events with intensities exceeding 100mm per hour occurring frequently in the summer, the risk of rainfall-induced landslides and debris flows is very high. In the aftermath of two destructive landslides in 1972, the Geotechnical Control Office set up the Slope Safety System. The system was designed to develop multiple initiatives, such as early warning mechanisms, community participation, and enforcement of geotechnical standards, to reduce landslide risk in a holistic way. The Slope Safety System has reduced the annual fatalities due to landslides by over 50% and there have been no fatalities in almost a decade (Malone, 1997).

Landslide susceptibility refers to the likelihood of landslide occurrence in a given study area, and is often evaluated by subdividing the area into zones based on factors influencing slope stability (Shano *et al.*, 2020). Throughout this work, we use the terms susceptibility and occurrence

probability interchangeably. A wide range of techniques has been used throughout the literature in the evaluation of landslide susceptibility, broadly classified into qualitative and quantitative approaches. Qualitative approaches, such as landslide inventory (Reichenbach *et al.*, 2018) and terrain analysis, rely on expert (heuristic) evaluation techniques, which are based on the evaluator's knowledge and experience (Corominas *et al.*, 2014). Quantitative approaches include statistical, deterministic, and probabilistic methods. Evolving from simple limit equilibrium models (Montgomery and Dietrich, 1994) to increasingly complex statistical and deep learning approaches, the expansion in the complexity and number of techniques has been matched by an ever-increasing number of publications in the scientific literature (Dong *et al.*, 2023).

This development of data-driven tools or landslide analysis has progressively addressed a range of themes, driven by research focused on performance-oriented solutions such as machine (Goetz *et al.*, 2015) and deep (Azarafza *et al.*, 2021) learning. In parallel, significant advances have emerged from targeted studies on: i) uncertainty estimation (Tanyas and Lombardo, 2020), sampling strategies for ii) stable (Steger *et al.*, 2016) and iii) unstable (Chang *et al.*, 2023) slopes, iv) space-time extensions (Lombardo *et al.*, 2020), as well as v) bias capture and removal (Steger *et al.*, 2021), and vi) variable selection (Budimir *et al.*, 2015) among others. Despite these advances, relatively little research has focused on adapting multivariate frameworks for hazard modelling, as opposed to susceptibility analysis.

The intensity of a single landslide—reflecting its potential impact—can be expressed in terms of its velocity, (Pudasaini and Krautblatter, 2021), kinetic energy (Chang *et al.*, 2017) or force (Tang *et al.*, 2014). However, whenever the scale of the study involves large geographic regions, such intensity-related measures are impossible to retrieve due to the costs involved with the geotechnical data acquisition (Van den Bout *et al.*, 2021). For this reason, few recent alternatives have been proposed to express landslide intensity over entire landscapes. Specifically, Lari *et al.* (2014) have proposed an intensity measure for rock falls under the assumption that failure sources are distributed according to a Poisson exponential family. Similarly, Lombardo *et al.* (2018) proposed a doubly stochastic structure (via a Log-Gaussian Cox Process - LGCP) to model the expected number of landslides per mapping unit. Note that from a purely statistical point of view, the LGCP tries to model a random intensity function that generates the observed point pattern. It is important here to differentiate the intensity function of the LGCP from the landslide intensity (a function of magnitude and velocity). Other alternatives include Taylor *et al.* (2018) who explored the possibility of expressing intensity as a function of landslide length-to-width ratios as a proxy indicator of the run-out characteristics. More

recently, Lombardo *et al.* (2021) presented a data-driven solution to the landslide intensity estimation by fitting a log-Gaussian model to global landslide planimetric areal data. This is where this thesis starts to build on the literature; specifically, in Chapter 2 we propose a Hurdle model to combine the probability of landslide occurrence with the expected planimetric extent. To our knowledge, this is the first joint probability model designed to statistically capture these two components of landslide hazard.

In this broad context, it becomes evident that no single solution exists to predict landslides and associated characteristics. One of the main determinants of the choice among these solutions has to do with data availability. In fact, most landslide inventories are expressed as point locations, these being diagnostic of where the failure mechanism occurred (Martha *et al.*, 2013). Therefore, the information on the planimetric area is not included as part of the inventory metadata. This is not the case for landslide polygonal inventories (Guzzetti *et al.*, 2012), which are more complex to define. Notably, deep-learning automated solutions have improved the situation in recent years, especially in addressing the need for multi-temporal landslide mapping (Bhuyan *et al.*, 2023).

An exploration into a scenario where we had available the point locations and the polygonal landslide inventory is outlined in Chapter 3. With both data types available, we propose a model comparison between the original model (heuristic), a susceptibility model (areal), and a point process model (point/geostatistical) for debris flows occurring across Scotland. This work is part of a project in defining a new data-driven landslide susceptibility model for the British Geological Survey (BGS).

On a whole, this work aims to integrate geostatistical and point-based data with environmental variables, often defined over geomorphologically meaningful discretisations of the study area, to ensure that landslide processes are represented in a geologically coherent manner. This thesis contributes to the analysis of such data–often at regional to large scale–by employing modelling approaches commonly used in environmental applications, while incorporating statistical techniques tailored to capture the complexity of landslide phenomena.. A central objective is to present model outputs in a format that is interpretable and useful for geoscientists, with a strong emphasis on preserving the intrinsic structure and resolution of the original data. This is specifically covered in Chapter 4.

While exploring the way high-resolution data could be implemented in the modelling archetype, the question of sufficient representation arose. In particular, that of dynamic environmental covariates such as precipitation and temperature. In the models explored in Chapters 2, 3 and 4, these covariates were reduced to a summary statistic such as mean annual precipitation. This prompted an exploration into statistical models that would preserve the functional nature of these dynamic

variables, specifically precipitation, and is described in Chapter 5.

This research aims to provide a comprehensive statistical assessment of landslide hazard, integrating the three fundamental components in a way that has not been previously achieved using statistical methods. The motivation behind this work was to develop and present a robust statistical framework for landslide hazard that could serve both methodological and practical purposes. We sought to employ a variety of statistical models to explore how landslides can be represented, with the dual objective of identifying potentially optimal approaches and improving upon existing statistical applications and implementations encountered. This motivation is particularly evident in Chapter 5, where we introduced a space–time stochastic partial differential equation (SPDE; Lindgren *et al.* 2011) smoothing term within the Mixed Generalised Additive Model Computation Vehicle with Automatic Smoothness Estimation package (`mgcv`), following the limitations we faced with modelling functional data using the integrated nested Laplace approximation (INLA; Rue *et al.* 2009). The timing of this research is especially pertinent given the creation and popularisation of the INLA methodology and the SPDE approach - developments that were simply not available in previous years, when such analyses would have been computationally prohibitive. This context is further strengthened by the work presented in Chapter 3, in which the BGS enhanced the National Landslide Database using satellite imagery to support a data-driven landslide susceptibility model. This advancement replaced a heuristic standard which had remained unchanged for two decades. The remainder of this introduction highlights a brief literature review with a focus on the geomorphological background, key statistical concepts, and modelling frameworks that are used throughout the rest of the thesis. Where necessary, these ideas are expressed more precisely as needed.

1.1 Geomorphological background

Reichenbach *et al.* (2018) critically review an extensive database of 565 peer-reviewed articles from 1983 (note that the first key paper on the instability of slopes due to the geological structure was written by Neuland (1976)) to 2016 that include statistically-based landslide susceptibility models. The review finds that there is a significant heterogeneity of thematic data types, modelling approaches (logistic regression, neural network analysis and data overlay) and model assessment and prediction, with only a handful of studies assessing the model uncertainty. Their recommendations can be summarised as follows: reduce the existing geographical bias by concentrating efforts in landslide susceptibility modelling in areas that have not been investigated before, by using multiple and complementary

mapping techniques, use geomorphological and multi-temporal inventories to test the predictive performance of the models, select an appropriate mapping unit with full knowledge of advantages and limitations of selection, to combine multiple model attempts to create an 'optimal' susceptibility zone, thus reducing model errors; and finally, to fully investigate the performance of the susceptibility model by taking into account model errors, distribution of model outcomes and model uncertainties. The review concludes by outlining future challenges that could arise within the literature, the most notable perhaps being that the susceptibility is not, as previously assumed, stationary (Samia *et al.*, 2017) and that the period of validity for susceptibility models may be in question. Reichenbach *et al.* (2018) can therefore be considered a guide on how to construct an investigation into landslide susceptibility by critical evaluation.

Expanding on the idea that susceptibility may not be stationary, Samia *et al.* (2017) aims to quantify landslide path dependency by using a follow-up landslide fraction, which incorporates the number of follow-up landslides and total landslides in a younger period, the area of interest and the area affected by landslides in an older period. They show that path dependency has an exponential response over time. In the study area, Collazzone, a municipality in central Italy, susceptibility increases by a factor of 15 immediately after an earlier landslide and decreases back to pre-landslide values in approximately 25 years. Landslide size and shape play a role in the probability of follow-up landslides, with larger and rounder landslides having a higher chance. Other geological covariates are seen to have an effect on predicting follow-up landslides, an example being topographic wetness index (TWI)—areas that were found to be wetter had a higher slope instability and hence positively influenced the chance of a follow-up landslide. A logistic regression model to predict follow-up landslide occurrence based on geographical/geological covariates predicted landslide occurrence with 60% confidence. How the output of this model might affect landslide susceptibility moving forward in statistical landslide hazard modelling is yet to be determined.

Finding data on landslides can be challenging, with inventory options including areal maps, landslide catalogues, the media/news, etc. Kirschbaum *et al.* (2015) present a new available global landslide catalogue (GLC) that provides a foundation for the evaluation of spatial and temporal trends in landslides from 2007 to 2013. The GLC considers mass movements triggered by rainfall which have been recorded in disaster databases, reports, the media or other sources. This new catalogue utilises a satellite-based precipitation record to identify patterns in rainfall, extreme enough to be aligned with landslide reports. The study found a clear relationship between the majority of the GLC points and extreme rainfall (classified as greater than the 90th percentile) and that the GLC could be used

to define new extreme rainfall thresholds, giving the ability to locate areas where perhaps landslides have been under-reported. The precipitation analysis indicated that extreme rainfall - the type to trigger landslides - may be better monitored by using a daily threshold, and that sampling should be taken in a very small window of accumulation to account for areas where the weather changed dramatically.

Having reviewed the construction of landslide hazard models and data inventories, we now turn to modelling techniques. Guzzetti *et al.* (2005) proposes a probabilistic model applied in the Staffora River basin, Italy, to predict the location, frequency and size of landslides. The paper has worked through the probability of landslide size (frequency-area distribution), the probability of landslide occurrence (mean recurrence and Poisson probability model) and the spatial probability of the slope failures and so, after assuming independence, these probabilities were multiplied together to achieve landslide hazard (the probability that a mapping unit will be affected by landslides that are greater than a certain size, during a certain time and because of certain environmental variables). The probability of the landslide size was calculated using a multi-temporal inventory map spanning 45 years where 2390 landslide were recorded. The inventory map paved the way for the estimated frequency of landslide occurrence in each slope unit over the 45 years to be calculated. Once the estimated landslide frequency was calculated, Guzzetti *et al.* (2005) used this alongside the rate of landslide occurrence over the time period for each slope unit (SU, hereafter) and the landslide recurrence (the expected time between each subsequent failure) for each SU to estimate the exceedance probability by using a Poisson probability model. The exceedance probability of having more than one landslide in each SU relies on the assumption that the underlying failure process remains stationary.

Although landslide hazard has been defined by occurrence probability and explained by environmental variables, there is no account for any spatial dependency - any pattern in the land that may influence the instability of the slopes. Lombardo *et al.* (2018) aims to account for the latent spatial effects in landslide modelling, specifically when looking at multiple debris flow scenarios from a precipitation trigger - as was the case in Messina, Sicily. This is done by treating each landslide triggering location as part of a random point pattern, which is different to the usual presence-absence structure used by Guzzetti *et al.* (2005). Lombardo *et al.* (2018) use a Poisson probability distribution for event counts in small-area slope units, thus extending the binary scenario to the effect that the distribution characterises the random number of events in a given unit, which can be modelled continuously over space. The paper investigates four possible models of the log-Gaussian Cox Process (LGCP) type, see Section 1.6, with a latent spatial effect and the inclusion of linear or non-linear covariates. These

four models are compared to a baseline model following the logistic regression approach, which helps to highlight the differences and advantages of the Poisson process approach. The model with the latent spatial effect demonstrated superior predictive performance.

In Lombardo *et al.* (2019b), the authors provide a detailed tutorial on fitting a Cox point process model for landslide spatial prediction using the integrated nested Laplace approximation (INLA; Rue *et al.* 2009). This work follows the novel approach described in Lombardo *et al.* (2018) and outlines three key steps for preparing and constructing the dataset: (i) compiling a point database of landslide-triggering locations, (ii) selecting a relevant set of covariates to explain landslide distribution, and (iii) defining appropiate mapping units for model implementation. The tutorial also covers data pre-processing, including the standardisation and grouping of covariates, aggregation of landslide counts at the mapping unit level, and construction of an adjacency graph structure to fit a Besag model (Besag *et al.*, 1991) for the latent spatial effect. Throughout this research, the treatment of latent spatial effects is adapted to suit different data structures. These methodological variations are further discussion in Section 1.5 and elaborated upon within each project presented in this thesis.

Expanding on Lombardo *et al.* (2018), Lombardo *et al.* (2019a) proposes a Poisson point process for earthquake-induced landslide susceptibility, noting that the Poisson distribution is consistent over all spatial resolutions. A log Gaussian Cox process was used to model landslide patterns caused by the seismic shaking of the Lushan (2013) and Wenchuan (2008) earthquakes. From the estimated landslide intensity, landslide susceptibility was calculated. The analysis found that the latent spatial effect captured the spatial distribution of a landslide trigger and was a good proxy for ground-shaking information.

Lombardo *et al.* (2021) extends the traditional approach towards landslide susceptibility (Bernoulli distribution) and the approach to model landslide intensity (Lombardo *et al.*, 2018) by modelling the size of landslides in a given mapping unit using a specific class of models - namely, latent Gaussian models, see Section 1.3. This was implemented using the log-Gaussian probability distribution to explain which characteristics of earthquake-induced landslides influence the largest landslide in the mapping unit and which influence the sum of the landslide areas in the mapping unit. The models were applied to a global dataset of earthquake-induced landslides, and spatial information was added to the modelling procedure in the same manner as previous studies, although the latent spatial effect was excluded. The reason behind this choice came from the large geographical extent and distribution of landslides (see Figure 1 in Lombardo *et al.* (2021)). Since the lack of spatial effects when modelling data with a strong spatial component will likely underestimate the uncertainty, the authors rely on a

Bootstrap procedure to assess the uncertainty of parameter estimates.

1.2 Bayesian statistics

Bayesian statistics is named after the 18$^{\text{th}}$-century minister Thomas Bayes, who, in 1763, presented a paper that was marked the first use of concepts now considered foundational to the Bayesian framework (O'Hagan, 2004). However, it was not until the 1950s that Bayesian methods began to gain broader acceptance and application. Bayesian statistics is a data analysis approach grounded in Bayes' theorem. Unlike frequentist methods, it utilises prior information—whether derived from existing data or based on expert knowledge—to estimate quantities of interest. These quantities, typically model parameters denoted by $\boldsymbol{\theta}$, are treated as random variables. The goal is to infer their posterior distribution, $\pi(\boldsymbol{\theta} \mid \boldsymbol{y})$, given the observed data $\boldsymbol{y}$. Bayes' theorem expresses this as:

$$\pi(\boldsymbol{\theta} \mid \boldsymbol{y}) = \frac{\pi(\boldsymbol{y} \mid \boldsymbol{\theta})\pi(\boldsymbol{\theta})}{\pi(\boldsymbol{y})}$$

where $\pi(\boldsymbol{y} \mid \boldsymbol{\theta})$ is the likelihood of the data, $\pi(\boldsymbol{\theta})$ is the prior distribution of the model parameters and $\pi(\boldsymbol{y})$ is the marginal likelihood. The likelihood function represents the probability of observing the data given specific parameter values, and the priors represent the 'best guess' for the true values of the model parameters as probability distributions. However, the marginal likelihood $\pi(\boldsymbol{y})$ is often computationally intractable. As a result, the posterior is typically expressed up to a proportionality constant by focusing on the numerator:

$$\pi(\boldsymbol{\theta} \mid \boldsymbol{y}) \propto \pi(\boldsymbol{y} \mid \boldsymbol{\theta})\pi(\boldsymbol{\theta}).$$

The main distinction between frequentist and Bayesian methods is that the Bayesian approach uses prior information. This is a strength and also a source of controversy. Prior information is, in principle, subjective. If my prior information differs from yours then we will ultimately have different posterior distributions. This raises the question of whether science can ever be entirely objective. Additionally, in the aim to reflect the nature of the data as closely as possible, subjectivity in the prior distribution is minimised by basing the prior information on factual and expert reasoning and through the accumulation of data, possible variability in the prior can be resolved (O'Hagan, 2004). Through Bayes' theorem, the prior information is weighted. If the prior is unclear or weak, then it will get a negligible weight and the posterior will be based solely on the data. The converse is likewise true.

One of the main advantages of the Bayesian approach is that, by incorporating prior information—even if subjective—it often more accurately reflects the nature of the problem. Additionally, the computation tools now available for Bayesian statistics allow for more complex problems to be tackled than in a frequentist setting. Lastly, the idea of 'uncertainty' is wholly covered in a Bayesian approach. 'Uncertainty' may refer to inherent randomness or to incomplete knowledge. Bayesian methods address both, whereas frequentist approaches typically account only for the former (O'Hagan, 2004).

1.3 Latent Gaussian Models

Bayesian inference is particularly suitable for a subclass of structured additive regression models known as latent Gaussian models (LGMs). For LGMs, the response is assumed to belong to the exponential family with a flexible prior structure that allows for localised dependencies and quantification of uncertainty (Sigrist, 2022). Throughout this thesis, we assume that observations are conditionally independent given a latent field and a set of hyper-parameters (Blangiardo *et al.*, 2013). For instance, if $\boldsymbol{y}$ is the vector of observations with components $y(\mathbf{s}), \mathbf{s} \in \mathcal{S}$, where $\mathcal{S}$ is the study region, $\boldsymbol{x}$ with components $x(\mathbf{s}), \mathbf{s} \in \mathcal{S}$, is a latent Gaussian random field (GRF) and $\boldsymbol{\theta}$ is the vector of hyper-parameters, a latent Gaussian field can be hierarchically represented as follows:

$$\begin{aligned} \boldsymbol{y} \mid \boldsymbol{x}, \boldsymbol{\theta}_1 &\sim \prod_{\mathbf{s} \in \mathcal{S}} \phi(y(\mathbf{s}) \mid x(\mathbf{s}), \boldsymbol{\theta}_1) \\ \boldsymbol{x} \mid \boldsymbol{\theta}_2 &\sim \mathcal{N}(\boldsymbol{\mu}_{\boldsymbol{x}}(\boldsymbol{\theta}_2), \boldsymbol{Q}_{\boldsymbol{x}}^{-1}(\boldsymbol{\theta}_2)) \\ \boldsymbol{\theta} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)^\top &\sim \pi(\boldsymbol{\theta}). \end{aligned} \tag{1.1}$$

where $\phi(\cdot)$ is the specified probability mass function or density characterising the observations. The latent Gaussian random field $\boldsymbol{x}$ describes the underlying dependence and trends in the data, which can be on a spatial or spatiotemporal scale. For a finite number of locations, the GRF is characterised by a mean vector $\boldsymbol{\mu}_{\boldsymbol{x}}(\boldsymbol{\theta}_2)$ and precision matrix $\boldsymbol{Q}_{\boldsymbol{x}}(\boldsymbol{\theta}_2)$. The vector of hyper-parameters $\boldsymbol{\theta}_2$ accounts for the variability and the strength of dependence of the random field (Blangiardo *et al.*, 2013), while $\boldsymbol{\theta}_1$ accounts for additional hyper-parameters in the likelihood.

The GRF is characterised by a flexible linear predictor with an additive structure that accommodates spatial and temporal descriptors. In a spatiotemporal context, this linear predictor can be

expressed as:

$$\eta(\mathbf{s},t) = \alpha + \sum_{m=1}^{M} \beta_m w_m(\mathbf{s},t) + \sum_{k=1}^{K} f_k(z_k(\mathbf{s},t)) + u(\mathbf{s},t) \tag{1.2}$$

Above, $\mathbf{s} \in \mathcal{S}$ (the study region), $t \in T$ (some set of time points), α is an intercept, and $(w_1(\mathbf{s}), ..., w_M(\mathbf{s}))^T$ is a subset of covariates available in the data with linear coefficients $\boldsymbol{\beta} = (\beta_1, ..., \beta_M)^T$. The functions $f = \{f_k(\cdot), ..., f_K(\cdot)\}$ are non-linear effects defined in terms of covariates $(z_1(\mathbf{s}), ..., z_K(\mathbf{s}))^T$. The specific form of the functions $f_k(\cdot)$ in this work is that of a Gaussian random walk of order 1 (RW1) or a Gaussian random walk of order 2 (RW2). The RW1 for a Gaussian vector $\boldsymbol{x} = (x_1, ..., x_n)^\top$ is constructed assuming independent Gaussian increments (Krainski *et al.*, 2018). Specifically,

$$\Delta x_i = x_i - x_{i-1} \sim \mathcal{N}(0, \tau^{-1}),$$

where the index i is the position of the covariate associated with y_i used to define the model under the assumption that covariate values are ordered increasingly. The density for $\boldsymbol{x}$ is derived from its $n-1$ increments as

$$\pi(\boldsymbol{x} \mid \tau) \propto \tau^{\frac{n-1}{2}} \exp\left\{\frac{-\tau}{2}\sum(\Delta x_i)^2\right\} = \tau^{\frac{n-1}{2}} \exp\left\{\frac{-1}{2}\boldsymbol{x}^T\boldsymbol{Q}\boldsymbol{x}\right\}$$

where $\boldsymbol{Q} = \tau\boldsymbol{R}$ and $\boldsymbol{R}$ is a matrix reflecting the neighbouring structure.

The term $u(\mathbf{s},t)$ is the spatiotemporal random effect, typically characterised by a Gaussian process with a stationary or non-stationary Matérn covariance structure (Matérn, 1986b) evolving in time according to some pre-specified structure, typically following an autoregressive structure. In general, the covariance can be a function of the distance between these locations only or it can be a location-dependent function that varies across space (see Section 1.5). In a purely spatial application, the notation t can be dropped from the terms in the linear predictor, and the spatial random effect is simply $u(\mathbf{s})$.

1.4 The Integrated Nested Laplace Approximation

LGMs offer a flexible framework for modelling spatial and spatio-temporal phenomena, as they can effectively capture local correlation structures and quantify associated uncertainty Sigrist (2022). As seen in the linear predictor, covariates and smooth effects can be incorporated in a linear or non-linear way. However, this introduces a strong dependency between the latent field and the model's hyper-parameters Rue *et al.* (2009), which can complicate inference methods. Markov Chain Monte

Carlo (MCMC) methods are a commonly used inference technique, but the strong dependency often means slow convergence and high computational costs. To combat this, Gaussian approximations can be introduced. Rue *et al.* (2009) reviews two approximation methods: variational Bayes (VB) and expectation-propagation (EP). These approaches ultimately have iterative solutions which, although designed to be faster than MCMC, often do not result in significant computational savings. Another approximation approach is the integrated nested Laplace approximation (INLA; Rue *et al.* 2009), which is now described in detail.

From the posterior marginals, $\pi(x_i \mid \mathbf{y}), i = 1, ..., n$, approximate conditional densities can be used to construct nested approximations. The approximate marginal posterior of hyper-parameters $\boldsymbol{\theta}$ is noted as

$$\tilde{\pi}(\boldsymbol{\theta} \mid \boldsymbol{y}) \propto \left. \frac{\pi(\boldsymbol{x}, \boldsymbol{\theta}, \boldsymbol{y})}{\pi_G(\boldsymbol{x} \mid \boldsymbol{\theta}, \boldsymbol{y})} \right|_{\boldsymbol{x}=\boldsymbol{\mu}(\boldsymbol{\theta})} \tag{1.3}$$

where $\pi_G(\boldsymbol{x} \mid \boldsymbol{\theta}, \boldsymbol{y})$ is the Gaussian approximation to the full conditional of $\boldsymbol{x}$, $\pi(\boldsymbol{x} \mid \boldsymbol{\theta}, \boldsymbol{y})$, and is found from a second-order expansion of the likelihood around its mode, $\boldsymbol{\mu}(\boldsymbol{\theta})$ (Rue *et al.*, 2009). The posterior marginals for $\boldsymbol{\theta}$ can then be approximated from 1.3 by integration with respect to $\boldsymbol{\theta}$. Marginal conditional posteriors of the elements of $\boldsymbol{x}$ are found from the joint Gaussian approximation

$$x_j \mid \boldsymbol{\theta}, \boldsymbol{y} \sim \mathcal{N}((\boldsymbol{\mu}(\boldsymbol{\theta}))_j, (\boldsymbol{\Omega}_{\boldsymbol{x}}^{-1}(\boldsymbol{\theta}))_{jj})$$

where $\boldsymbol{\theta}$ needs to be integrated out to find the marginal posterior of x_j. Specifically,

$$\tilde{\pi}(x_j \mid y) = \int \pi_G(x_j \mid \boldsymbol{\theta}, \boldsymbol{y}) d\boldsymbol{\theta} \approx \sum_{k=1}^{K} \pi_G(x_j \mid \boldsymbol{\theta_k}, \boldsymbol{y}) \tilde{\pi}(\boldsymbol{\theta_k} \mid \boldsymbol{y}) \delta_k \tag{1.4}$$

where $\boldsymbol{\theta_k}$ are K integration points and δ_k are corresponding area weights (Van Niekerk *et al.*, 2023).

By using the Gaussian marginal derived from $\tilde{\pi}_G(\boldsymbol{x} \mid \boldsymbol{\theta}, \boldsymbol{y})$, the density of $x_i \mid \boldsymbol{\theta}, \boldsymbol{y}$ can be approximated and numerically integrated with respect to the hyper-parameters, and thus the posterior marginals of the latent field can be obtained. Here, difficulties in controlling the error of the approximations occurred because the points for integration could not be chosen in an automatic and adaptive way. Rue *et al.* (2009) addresses these issues by reapplying the Laplace approximation to the conditional posterior marginals of the latent field, and proposes a faster correction to the Gaussian approximation $\tilde{\pi}(x_i \mid \boldsymbol{\theta}, \boldsymbol{y})$.

The Laplace approximation should only be applied to near-Gaussian densities (Rue *et al.*, 2017). In the case of having the product of a Gaussian density and a non-Gaussian density, Rue *et al.* (2017) conditions on some well-behaved function that writes the conditional marginal posterior of this function as an integral of densities that are almost Gaussian. Here, this non-Gaussian density is

conditioned into a correction of a Gaussian prior, which, when computed, is close to Gaussian. Although this substitution increases complexity—since numerical integration must be performed across multiple values, the same concept can be applied to LGMs. Therefore, LGMs consist of a vector of hyperparameters and a latent Gaussian random field that encodes the model's structure. The Laplace approximation requires a Gaussian approximation of the denominator, so the hyperparameters have a condition that is approximated as a Gaussian Markov Random Field (GMRF), but the condition on the latent field is slightly more challenging. Integration over the conditioned hyperparameters would incur an exponential computational cost from the dimension of the required integration and the latent field needs to be approximated in a subset of its range $i = 1, ..., n$, where n can be extremely large. To overcome this, a central composite design (Box and Wilson, 1951) is employed to balance efficiency and computational costs. This is done by placing the integration points on a sphere, suitably located for the joint posterior of the hyperparameters, and equally weighting the points to determine the relative weight to the central point given a Gaussian posterior. To approximate the latent field n times, a Taylor expansion around the mode of the Laplace approximation is used. This way, the posterior marginals for the latent field are approximated by a mixture of skew-Normal distributions.

In the classical formulation, the marginal posteriors of the linear predictors are calculated within the marginal posterior distributions of the latent field. However, in the recent update by Van Niekerk *et al.* (2023), a low-rank variational Bayes correction to the mode can be calculated to solve for the posterior distributions of the linear predictors.

1.5 The Stochastic Partial Differential Equation approach

Gaussian fields were previously introduced in general terms. In this section, we explore in more detail this and dissect the idea of the stochastic partial differential equation approach in the context of spatial statistics. Let $\mathbf{s}$ be any location in a study area $\mathcal{D}$ and let $u(\mathbf{s})$ be the random spatial effect at that location (like the term in 1.2 if t is dropped). As mentioned in Section 1.3, $u(\mathbf{s})$ is a stochastic process, meaning it varies continuously in a random way over space. If $u(\mathbf{s})$ is assumed to be Gaussian, then it is said to be a continuously indexed Gaussian field (GF). This implies that it is possible to collect these data at any finite set of locations within the study region. The process $u(\mathbf{s})$ is usually assumed isotropic and stationary, meaning that there is a constant mean and a covariance that depends only on the distance between points and not on their relative positions. Lindgren *et al.* (2011) represent a GF with Matérn covariance (Matérn, 1986b) as a Gaussian Markov Random Field

(GMRF). The GMRF provides a sparse representation of the GF through a sparse precision matrix. This representation is done through a solution of the SPDE using the finite element method.

Before presenting the solution, we first review the Matérn covariance function—one of the most widely used families in spatial statistics due to its flexible local behaviour and ability to model varying degrees of smoothness (Guttorp and Gneiting, 2006a).. It is defined as

$$R_M(h) = \frac{\sigma^2}{2^{\nu-1}\Gamma(\nu)}(\kappa h)^\nu K_\nu(\kappa h),$$

where ν is a positive smoothness parameter, κ is a positive scale parameter, σ^2 is the variance of the corresponding random field, K_ν is the Bessel function of the second kind of order ν and Γ is the Gamma function. This formulation controls both the differentiability of the process (via ν) and the correlation range (via κ). As stated in Lindgren *et al.* (2011), a Gaussian Matérn field solves the SPDE

$$(\kappa^2 - \Delta)^{\frac{\alpha}{2}} u = \mathcal{W},$$

where Δ is the Laplacian, $\mathcal{W}$ is Gaussian white noise and $\alpha = \nu + d/2$. Lindgren *et al.* (2011) shows that as the smoothness parameter ν increases, the precision matrix in the GMRF representation becomes less sparse. They also demonstrate how the method applies to irregular spatial grids, which are typical in real-world applications.

The SPDE approach approximates this covariance structure by approximating the precision matrix and finding an SPDE whose solution has the desired covariance structure. Lindgren *et al.* (2011) show how to find the solution by discretising the study region with a triangulation mesh and representing the stochastic process as a sum of basis functions multiplied by a set of coefficients that form a GMRF for which there are already methods for fast computation and for which the dependencies are defined by the triangulation of the spatial domain. The construction of a mesh is vital in representing the spatial process and affects the speed and accuracy of the estimation. A sensitivity analysis for how the required estimates vary over different mesh constructions is often conducted. A Delaunay triangulation is typically used, as it maximises the minimum interior angle, smoothing the transitions between small and large triangles. Extra vertices are added heuristically to try to minimise the total number of triangles needed to meet shape and size constraints.

The finite element representation of the SPDE solution (Brenner, 2008) can be written as

$$u(\mathbf{s}) = \sum_{k=1}^{K} \psi_k(\mathbf{s}) w_k \tag{1.5}$$

for chosen basis functions ψ_k and some Gaussian distributed weights w_k. The number of vertices in the triangulation is defined here by K. The associated weights determine the value of the field at each vertex, while values within the interior of the triangles are obtained by linear interpolation. To obtain a finite-dimensional solution, the distribution of the representation weights must satisfy the stochastic weak formulation of the SPDE (equation 7 in Lindgren *et al.* (2011)) for a given set of basis functions. Building on this, Lindgren *et al.* (2022) refine the SPDE approach by constructing precision matrices in a manner that ensures robustness to the choice of spatial neighbourhood-defining graphs.

1.6 Spatial Point Processes

Spatial point processes describe short-range interactions that explain the spatial configuration of observed points (Illian *et al.*, 2008). These points reflect the underlying processes driving the spatial pattern, including geophysical properties of the study area. A Poisson point process is a type of spatial process where the number of points within a given region follows a Poisson distribution (Baddeley *et al.*, 2007). A spatial Poisson process with a spatially varying random intensity $\lambda(s)$ is known as a Cox process (Møller *et al.*, 1998). This type of process is *doubly stochastic*, as it arises from an inhomogeneous Poisson process with a random intensity measure. When the logarithm of this random intensity is modelled as a Gaussian process, the resulting model is known as a log-Gaussian Cox process (LGCP; Bachl *et al.* (2019)). A key property of LGCPs is that the distribution is completely characterised by the intensity measure, making parameter estimation easy to interpret Møller *et al.* (1998). The random intensity function for a Cox process can be expressed as

$$\Lambda(\mathbf{s}) = \exp(\eta(\mathbf{s})), \tag{1.6}$$

and taking the logarithm yields the LGCP formulation:

$$\log(\Lambda(\mathbf{s})) = \eta(\mathbf{s}). \tag{1.7}$$

Here, $\eta(\mathbf{s})$ corresponds to the linear predictor defined in (1.2), with time t omitted.

As the class of LGMs include LGCPs, INLA can be used for fitting and Illian *et al.* (2012) developed a framework for this. By discretizing the observation window into $N = n_{row} \times n_{col}$ grid cells, s_{ij} where $i = 1, ..., n_{row}$ and $j = 1, ..., n_{col}$, the point pattern is then represented by the observed number of points, y_{ij}, in a grid cell. It is from the definition of an LGCP that the observed number of points in a grid cell can be considered as independent Poisson random variables $y_{ij} \sim \text{Poi}(\Lambda_{ij})$ where

$\Lambda_{ij} = \int_{s_{ij}} \lambda(s)\mathrm{d}s$ is the cell intensity. As the individual intensities for all grid cells are computationally infeasible to calculate, the approximation below is used

$$\Lambda_{ij} \approx \mid \mathbf{s}_{ij} \mid \exp(z_{ij}),$$

where z_{ij} is a value that represents the GRF $Z(\mathbf{s})$ within each cell s_{ij}, and $\mid \mathbf{s}_{ij} \mid$ is the area of each grid cell. If the formulation of grid cells is fine enough and the latent Gaussian field is appropriately defined, this approximation is useful (Waagepetersen, 2004) and converges to the true solution as the grid cell size tends to zero.

The role of the lattice (collection of grid cells) is to approximate the latent GRF and to approximate the locations of the points. As the quality of the likelihood approximation depends on the size of the grid (the finer the grid, the less information is lost), a very fine grid is desired for the approximation of the point locations. But having a very fine grid to approximate the GRF involves a high computational cost, and is not necessary. Therefore, Simpson *et al.* (2016) propose moving away from lattice-based methods towards a more flexible approach that can integrate local changes to the resolution of the likelihood approximation by having a finite-dimensional continuously specified GRF. The likelihood can be written as

$$\log \pi(y \mid Z) = \mid \Omega \mid - \int_{\Omega} exp\{Z(\mathbf{s})\} d\mathbf{s} + \sum_{i}^{n} Z(S_i),$$

which consists of a stochastic integral and the spatial field evaluated at the points of the data, which can be computed by using the SPDE models, while the integral can be approximated by a deterministic integration rule:

$$\int_{\Omega} f(s)ds \approx \sum_{i=1}^{p} \tilde{\alpha_i} f(\tilde{s_i}),$$

where the $\tilde{s_i}$'s are a set of fixed deterministic nodes and the $\tilde{\alpha_i}$ are associated fixed weights. The log-likelihood approximation can then be written as

$$\begin{aligned} \log \pi(y \mid z) &\approx C - \sum_{i=1}^{p} \tilde{\alpha_i} \exp\{\sum_{j=1}^{n} z_j \phi_j(\tilde{s_i})\} + \sum_{i=1}^{N} \sum_{j=1}^{n} z_j \phi_j(s_i) \\ &= C - \tilde{\alpha^T} \exp(A_1 z) + 1^T A_2 z, \end{aligned} \tag{1.8}$$

where C is a constant, $[A_1]_{ij} = \phi_j(\tilde{s_i})$ the matrix of values of the latent Gaussian field at the integration nodes, and $[A_2]_{ij} = \phi_j(s_i)$ is the evaluation of the latent Gaussian field at the observed points. Therefore, we have p integration points, used to approximate the log-likelihood integral term, n mesh nodes, and N number of observed points. The form of Equation 1.8 is Poisson in the sense that the approximate likelihood consists of $N + p$ independent Poisson random variables.

1.7 Data and software

The focus of my research lies in the statistical aspects of the models chosen and the theory implemented and extended. However, I was awarded the Jim Gatheral Travel Scholarship 2023, which allowed me to visit my external supervisor, Dr Luigi Lombardo, at the University of Twente. During this visit, the aim was to become acquainted with the process of data extraction necessary for my projects and consequently the refinement of the data.

I was introduced to the Geospatial Computing platform (Crib) for the Department, in which applications such as Geographic Resources Analysis Support System (GRASS), a geographic information system (GIS) and QGIS are located. I learned how to navigate GRASS GIS as this is important for the calculation of the resolution of the data. I was also taught how to use Google Earth Engine (GEE). From this, one can write JavaScripts to extract geographical/geological information from anywhere in the world, by defining the area of interest and specifying which mapping system to use (e.g. Shuttle Radar Topography Mission, Climate Hazards Group InfraRed Precipitation with Station data, etc.). One can then import this extracted information into GRASS GIS. Transforming the information from raster to vector and into the correct coordinate plane is also something I learned how to do. There exist scripts to pass into the console that first removes the flat areas in the area of study and then generate the slope unit (or some such desired resolution of data) in vector format, most commonly using an open-source software named `r.slopeunits` (Alvioli *et al.*, 2016), and one can also input the raster maps of the geographical information and extract the values of the spatial information at this desired resolution.

Alongside the physical data extraction software used for this research, the computing aspect and the main tool utilised is R, a programming language specifically designed for statistical computing, and RStudio, the integrated development environment for R. Implementing the inference technique, INLA, in R is convenient as it is implemented in the `R-INLA` library (Bivand *et al.*, 2015). Other key libraries that were used include `inlabru`, built for Bayesian spatial modelling using integrated nested Laplace approximations and `INLAspacetime` - a package to implement certain spatial and spatiotemporal models using the `cgeneric` interface in the INLA package. Lastly, the resulting hazard map production for model results was done through ArcGIS Pro.

The analyses presented in this work are available at: `https://github.com/BryceErin`. We can provide the data and code for Chapter 2 and 3, the code for the simulation studies presented in Chapter 4 and the code for the space-time SPDE smoother. For information on the Japanese

landslide data and the Chinese surface deformation data, contact the authors.

Chapter 2

A Hurdle model

This chapter presents a Hurdle model that jointly estimates the probability of landslide occurrence and its size, measured by planimetric extent. The approach produces a unified landslide hazard map and provides probabilistic estimates, along with uncertainty quantification, for exceedance probabilities at high extent levels. The work in this Chapter has been published as:

Bryce, E., Lombardo, L., Van Westen, C., Tanyas, H. and Castro-Camilo, D., 2022. Unified landslide hazard assessment using hurdle models: a case study in the Island of Dominica. *Stochastic Environmental Research and Risk Assessment*, 36(8), pp.2071-2084.

2.1 Introduction

Dominica is situated within the Atlantic hurricane belt and, as such, is highly vulnerable to high-intensity weather events, as evidenced by its long extreme-event history, stretching as far back as the Great Hurricane of 1780. In addition, the island's infrastructure and population are concentrated along the coastal areas, particularly in the south and west, where their situation and building conditions hardly help withstand one extreme natural hazard after another. In this chapter, the focus is on the aftermath of Hurricane Maria, which originated from a wave leaving the west coast of Africa on 12$^{\text{th}}$ September 2017. The wave moved westward across the Atlantic, creating deep convection, and consolidated into a tropical depression 580 nautical miles east of Barbados on the 16$^{\text{th}}$ September (Pasch *et al.*, 2018; Fobert *et al.*, 2021). By the evening of the 18$^{\text{th}}$, Maria had intensified and hit the Island of Dominica twelve hours later as a category five hurricane (Pasch *et al.*, 2018), causing a total of 9,960 landslides on the island, 8,576 of which were classified as debris slides (Van Westen and Zhang, 2018). The infrastructure and transport sectors were the main avenues of the USD 930.9 mil-

lion in damages incurred by Maria, with the agriculture and tourism sectors not far behind (ACAPS, 2018). This prompted a plea by Dominica's prime minister for the funds to make Dominica a fully climate-resilient nation, determined to protect their island (Gibbens, 2019).

The Government of the Commonwealth of Dominica, in conjunction with the Caribbean Disaster Emergency Management Agency and the Caribbean Development Bank, commissioned a post-disaster needs assessment from Hurricane Maria in order to estimate the total damage, the damages per sector, and to identify recovery needs (ACAPS, 2018). The reconstruction of the destroyed infrastructure was funded by a loan from the World Bank and was a part of the project "Enhancing Resilient Reconstruction in Dominica". The project promoted the idea that Dominica could limit the damage from natural hazards by improving the uptake of resilient building practices, aiming to accelerate short-term recovery and strengthen long-term resilience to climate-related risks. Landslide hazard assessment is critically valuable to this programme. It can help define land-use capability, detect areas where intervention is needed to stabilise slopes, and identify appropriate mitigation measures.

To help tackle the above challenges, this article builds a joint probability model to map the Island of Dominica by susceptibility and size of landslides. The model is able to detect unstable areas that could potentially host large landslides, which is key to identifying regions in need of hazard mitigation. The model belongs to the class of hurdle models, i.e., it is a two-part model that specifies a Bernoulli likelihood for landslide occurrence and a Gaussian likelihood for landslide log-sizes given a positive occurrence. The observations (presence/absence for the Bernoulli likelihood and log-sizes for the Gaussian likelihood) are assumed to be conditionally independent given a Gaussian latent process that drives the trends, dependencies and non-stationary patterns observed in the data. This latent process is therefore characterised by a linear predictor as defined in Section 1.3. This flexible additive structure is utilised to model the covariates' influence using fixed and random effects, which are also known as linear and non-linear effects in terms of their influence. The spatial dependence between locations is hereby characterised using a Gaussian Process with Matérn covariance structure. The Matérn family of covariance functions is widely used in spatial statistics due to its flexible local behaviour and interesting theoretical properties (Guttorp and Gneiting, 2006b; Stein, 2012). The Matérn covariance is represented accurately by the stochastic partial differential equation approach (SPDE; Lindgren *et al.*, 2011). Under a Bayesian framework, we assume relatively weak but highly interpretable Gaussian priors for all the model components and hyperparameters involved and fit our model using the integrated nested Laplace approximation (INLA; Rue *et al.*, 2009). INLA uses the Laplace approximation to compute posterior distributions of interest. It is developed for the vast

class of latent Gaussian models and implemented in the `R-INLA` library (Bivand *et al.*, 2015).

Throughout this chapter, the term susceptibility will be used interchangeably with the probability of occurrence (see Lombardo *et al.*, 2020; Titti *et al.*, 2021). As for the notion of hazard, this further extends the concept of susceptibility by looking at measurable quantities of hazardous processes, such as size and frequency (Guzzetti *et al.*, 1999). However, the literature on statistical applications in landslide science has never presented a model where the susceptibility is combined with the expected size of landslides, which is proposed here for the first time.

The remainder of this chapter is organised as follows. Section 2.2 introduces the variables considered for modelling. Section 2.3 outlines aspects of the methodology specific to this case study, with the basis and theory of much of the framework detailed in Chapter 1. In Section 2.4, the results of the model are summarised and a critical discussion provided in Section 2.5.

2.2 Data

After hurricane Maria, the Global Facility for Disaster Reduction and Recovery sponsored the CHARIM (Caribbean Handbook on Risk Information Management; link here) project, intending to build a comprehensive natural hazard assessment. As part of this assessment, a team from the University of Twente (NL) mapped a large-scale landslide inventory for the whole island (see Figure 2.1a), primarily using five scenes of Pléiades satellite images with a resolution of 0.5m, dated on the 23^{rd} September and 5^{th} October 2017, made available through UNITAR-UNOSAT (https://www.unitar.org/). Additionally, a series of Digital Globe Images were used after they were gathered for the Google Crisis Response. All the images were visually inspected by image interpretation experts. As a result, landslides were mapped as polygons, separating scarp, transport and accumulation areas, and classifying the landslides according to their types.

The inventory features 9960 landslides, with a cumulative planimetric area of 11.4 km^2, covering 1.5 % of the Island of Dominica. To aggregate the landslide information over space, a Slope Unit partition of Dominica was opted for (see Figure 2.1b). A Slope Unit (SU, hereafter) encompasses the geographic space between ridges and streamlines. In other words, they are half-basins of a given order (Amato *et al.*, 2019) extracted by maximising areas with homogeneous aspects and are particularly suited to support landslide models for they approximate the morpho-dynamic behaviour of landslides well (Carrara *et al.*, 1991). To partition the Island of Dominica into SUs, `r.slopeunits` was initially used, an open-source software developed by Alvioli *et al.* (2016). Later, this was refined

through manual editing in GIS to obtain a total of 3335 SUs.

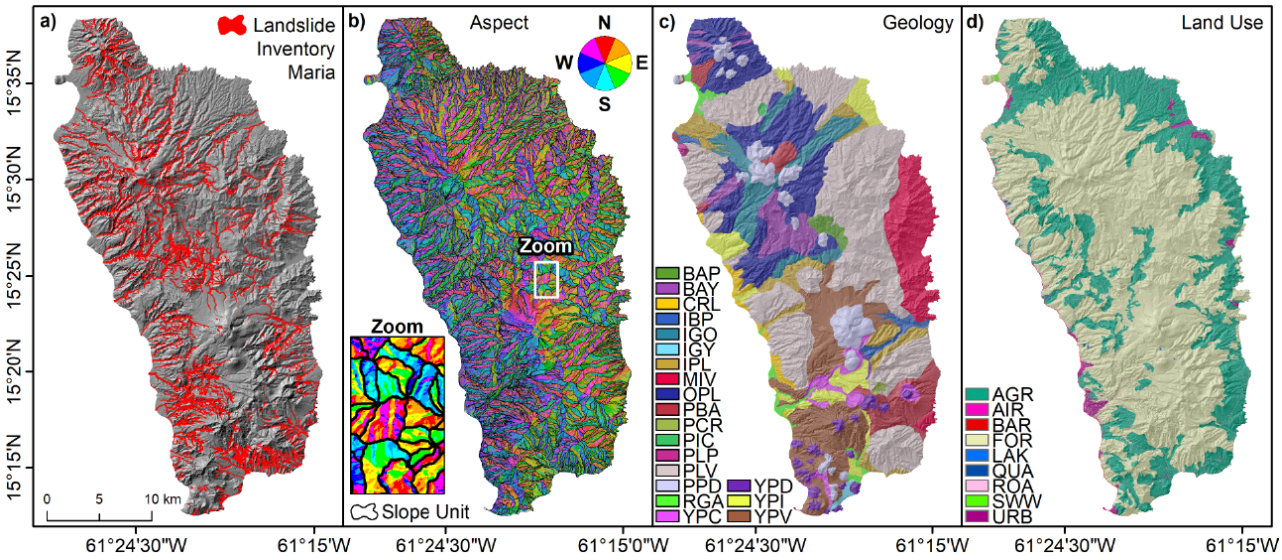

Figure 2.1: Summary of the data used in this work. Panel a) shows the landslide inventory over the shaded terrain. Panel b) shows the SU delineation over the aspect map. The area marked as "Zoom" in panel b) highlights the consistency of SU boundaries with changes in aspect. Panels c) and d) report the geological and land-use maps, respectively.

A two-step procedure was used to build the hurdle model at the SU scale. First, assigned to each SU was the sum of all planimetric landslide areas intersecting the SU itself. Then, to build a binary dataset, a landslide presence was assigned to SUs with positive landslide aggregated areas. Conversely, a landslide absence was assigned to SUs unaffected by slope failures. Finally, to build a dataset of landslide sizes per SU, SUs with positive landslide aggregated areas were extracted.

Covariates available for analysis are detailed in Table 2.1 They are a mixture of geographical and geological characteristics and constitute the morphology of Dominica's landscape. All the geographical variables are continuous, while the geological and land-use types are represented as proportions of each SU; for example, a SU could be 50% forest, 20% bare and 30% quarry.

Table 2.1: Summary of the Dominica dataset, including the responses and initial covariates' set.

Variable	Acronym	Type	Units
Landslide size	$A_L(\mathbf{s})$	Continuous response	m^2
Landslide occurrence	$O_L(\mathbf{s})$	Binary response	0=No, 1=Yes
SU area	SU_A	Continuous explanatory	m^2
SU perimeter	SU_P	Continuous explanatory	m
SU_P / SU_A	$SU_{P/A}$	Continuous explanatory	1/m
SU_P / $\sqrt{SU_A}$	$SU_{P/\sqrt{A}}$	Continuous explanatory	Unit-less
SU maximum distance	SU_{MD}	Continuous explanatory	m
SU maximum distance to area ratio	$SU_{MD/A}$	Continuous explanatory	1/m
SU_{MD} / $\sqrt{SU_A}$	$SU_{MD/\sqrt{A}}$	Continuous explanatory	Unit-less
Distance to nearest stream mean and SD	D2S μ and σ	Continuous explanatory	m
Eastness mean and SD	EN μ and σ	Continuous explanatory	Unit-less
Elevation mean and SD	EL μ and σ	Continuous explanatory	masl
Northness mean and SD	NN μ and σ	Continuous explanatory	Unit-less
Planform curvature mean and SD	PLC μ and σ	Continuous explanatory	1/m
Profile curvature mean and SD	PRC μ and σ	Continuous explanatory	1/m
Relative slope position mean and SD	RSP μ and σ	Continuous explanatory	Unit-less
Slope mean and SD	SLO μ and σ	Continuous explanatory	Degrees
Topographic wetness index mean and SD	TWI μ and σ	Continuous explanatory	Unit-less

2.3 Modelling framework

2.3.1 Hurdle model specification

As mentioned in Section 2.1, the aim is to build a hurdle model to detect areas of higher landslide susceptibility and sizes. The approach models the probability of observing a landslide in a SU using a Bernoulli distribution. Given that a landslide was observed, it also describes landslide log-sizes using a Gaussian distribution. It is assumed that the observations (presence/absence for the Bernoulli likelihood and log-sizes for the Gaussian likelihood) are conditionally independent given a latent Gaussian structure that drives the trends, dependencies and non-stationary patterns observed in the data. Both parts of the hurdle model belong to the class of latent Gaussian models.

To consider the models in terms of the latent Gaussian hierarchical representation;

$$\begin{aligned} \boldsymbol{y} \mid \boldsymbol{x}, \boldsymbol{\theta}_1 &\sim \prod_{\mathbf{s}\in\mathcal{S}} \phi(y(\mathbf{s}) \mid x(\mathbf{s}), \boldsymbol{\theta}_1) \\ \boldsymbol{x} \mid \boldsymbol{\theta}_2 &\sim \mathcal{N}(\boldsymbol{\mu}_{\boldsymbol{x}}(\boldsymbol{\theta}_2), \boldsymbol{Q}_{\boldsymbol{x}}^{-1}(\boldsymbol{\theta}_2)) \\ \boldsymbol{\theta} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)^\top &\sim \pi(\boldsymbol{\theta}). \end{aligned}$$

where $\phi(\cdot)$ is the Bernoulli probability mass function for landslide susceptibility or the Gaussian density for landslide log-sizes. The latent Gaussian random field $\boldsymbol{x}$ describes the underlying dependence and trends in the data. For a finite number of locations, the GRF is characterised by a mean vector $\boldsymbol{\mu}_{\boldsymbol{x}}(\boldsymbol{\theta_2})$ and precision matrix $\boldsymbol{Q}_{\boldsymbol{x}}(\boldsymbol{\theta_2})$. The vector of hyper-parameters $\boldsymbol{\theta_2}$ accounts for the variability and the strength of dependence of the random field (Blangiardo *et al.*, 2013), while $\boldsymbol{\theta_1}$ accounts for additional hyper-parameters in the likelihood. So, for instance, if $\phi(\cdot)$ is the Gaussian density, then $\boldsymbol{\theta_1}$ is equal to the Gaussian precision, which is the inverse of the variance.

Furthermore, this hierarchical representation is simplified into the form of a linear predictor; an additive sum of the model components, defined as

$$\begin{aligned} \eta_{\text{Bern}}(\mathbf{s}) &= \alpha + \sum_{m=1}^{M} \beta_m w_m(\mathbf{s}) + \sum_{k=1}^{K} f_k(z_k(\mathbf{s})) + u(\mathbf{s}) \\ \eta_{\text{Gauss}}(\mathbf{s}) &= \alpha + \sum_{m=1}^{M} \beta_m w_m(\mathbf{s}) + \sum_{k=1}^{K} f_k(z_k(\mathbf{s})) + u(\mathbf{s}) \end{aligned} \tag{2.1}$$

Above, $\mathbf{s} \in \mathcal{S}$ (the study region), α is an intercept, and $(w_1(\mathbf{s}), ..., w_M(\mathbf{s}))^T$ is a subset of covariates available in the data with linear coefficients $\boldsymbol{\beta} = (\beta_1, ..., \beta_M)^T$. The functions $f = \{f_k(\cdot), ..., f_K(\cdot)\}$ are non-linear effects defined in terms of covariates $(z_1(\mathbf{s}), ..., z_K(\mathbf{s}))^T$. The specific form of the functions $f_k(\cdot)$ in this work is that of a Gaussian random walk of order 1 (RW1). Finally, the term $u(\mathbf{s})$ is the spatio random effect, typically characterised by a Gaussian process with a stationary or non-stationary Matérn covariance structure (Matérn, 1986b). See Section 1.3 and 1.5 for more details.

Landslide susceptibility is modelled with a Bernoulli distribution. Specifically, $y(\mathbf{s}) = \text{O}_\text{L}(\mathbf{s}) \in \{0, 1\}$ and $\phi(y(\mathbf{s}) \mid \eta_{\text{Bern}}(\mathbf{s})) \equiv \text{Bern}(p(\mathbf{s}))$, where $p(\mathbf{s}) = \Pr\{\text{O}_\text{L}(\mathbf{s}) = 1\}$; see Figure 2.2(a) for the spatial distribution of $\text{O}_\text{L}(\mathbf{s})$. Note that there are no hyperparameters for the Bernoulli likelihood, i.e., the vector $\boldsymbol{\theta_1}$ is empty. The probability $p(\mathbf{s})$ is related to the linear predictor $\eta(\mathbf{s})$ through the logit link, so $p(\mathbf{s}) = \exp\{\eta_{\text{Bern}}(\mathbf{s})\}/(1 + \exp\{\eta_{\text{Bern}}(\mathbf{s})\})$.

The landslide size distribution is positively skewed, with extremely large values elongating the right tail. In cases like this, it is standard practice to use a monotonic transformation such as the natural logarithm to obtain a roughly symmetric, Gaussian-like distribution (Lombardo *et al.*, 2021).

Given that a landslide has occurred in the SU $\mathbf{s} \in \mathcal{S}$, the log-size is modelled using a Gaussian distribution such that; $y(\mathbf{s}) = \log\{A_L(\mathbf{s})\} \mid \mathrm{O_L}(\mathbf{s}) = 1$ and $\phi(y(\mathbf{s}) \mid \eta_{\mathrm{Gauss}}(\mathbf{s}), \boldsymbol{\theta}_1) \equiv \mathcal{N}(\eta_{\mathrm{Gauss}}(\mathbf{s}), \tau^{-1})$, where $\boldsymbol{\theta}_1 = \tau$. Figure 2.2(b) shows the spatial distribution of $\log\{A_L(\mathbf{s})\}$. The linear predictor is linked to the Gaussian mean via the identity link. The global precision hyperparameter τ (reciprocal of the standard deviation) determines the concentration of all values $y(\mathbf{s})$ around their mean $\eta(\mathbf{s})$, $\mathbf{s} \in \mathcal{S}$.

The linear predictors $\eta_{\mathrm{Bern}}(\mathbf{s})$ and $\eta_{\mathrm{Gauss}}(\mathbf{s})$ follow the general form as described in Equation 2.1. Nonetheless, their specific form depends on the influence of the covariates in Table 2.1 over the landslide susceptibilities and sizes. Therefore, variable selection is conducted using a stepwise forward procedure for most of the covariates in both parts of the hurdle model. This procedure was based on numerical techniques such as the Deviance Information Criterion and the Watanabe-Akaike information criterion (DIC and WAIC, respectively; Meyer, 2014, Gelman *et al.*, 2014), and graphical methods such as the probability integral transform (PIT; Gneiting *et al.*, 2007) and fitted versus observed plots. Due to their definition or interpretability, some covariates were not tested this way. Instead, they were included or excluded based on expert opinion. Specifically, expert advice was considered for land-use types and lithology types (included linearly as they are represented as a proportion of a SU and therefore, their sum is constrained to 1); mean Eastness and mean Northness (included linearly as they are complementary measurements representing the sine and cosine of the aspect of a SU); mean slope (included non-linearly based on previous analysis, see, e.g., Tanyaş *et al.*, 2022); and the standard deviations of all covariates (included linearly since non-linear standard deviations lack reasonable interpretation). All the covariates related to the perimeter were excluded due to collinearity issues. The selected covariates and the specific form they enter into the linear predictor equation for each model are detailed in Table 2.2.

2.3.2 Prior specification

Here the form of the third layer of the hierarchical representation in Equation 1.1 is described, i.e., the prior distributions for the likelihood hyperparameters $\boldsymbol{\theta}_1 = \tau$ (the Gaussian precision) and the hyperparameters of the model components in the linear predictor, $\boldsymbol{\theta}_2 = (\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\tau}_0, \boldsymbol{\sigma}^2, \boldsymbol{r})^\top$. Here, $\boldsymbol{\alpha}$ is a vector of length two that contains the intercept for both likelihood models, $\boldsymbol{\beta}$ is a vector of length 48 that contains the fixed effects for both likelihood models and $\boldsymbol{\tau}_0$ is a vector of length ten that contains the RW1 precision parameters for both likelihood models; see Table 2.2. The vectors $\boldsymbol{\sigma}^2$ and $\boldsymbol{r}$ have length two and contain, respectively, the marginal variances and range parameters of the

Table 2.2: Summary of selected covariates for both likelihood models.

Likelihood	Fixed effects	Random effects
Bernoulli	SU_{MD}, $SU_{MD/A}$, $SU_{MD/\sqrt{A}}$, D2Sμ, D2Sσ, ENμ, ENσ, NNμ, NNσ, ELσ, PLCμ, PLCσ, PRCσ, RSPμ, RSPσ, SLOσ, TWIμ, TWIσ, RGA, YPC, YPV, YPI, BAY, IGO, IPL, IGY, OPL, PBA, PPD.	SU_A, ELμ, SLOμ, PRCμ.
Gaussian	$SU_{MD/A}$, $SU_{MD/\sqrt{A}}$, D2Sσ, ENμ, ENσ, NNμ, NNσ, ELσ, PLCμ, PLCσ, PRCμ, PRCσ, RSPμ, RSPσ, SLOσ, TWIσ, YPV, BAY, OPL.	SU_A, SU_{MD}, D2Sμ, ELμ, SLOμ, TWIμ.

Matérn covariance of $u(\mathbf{s})$ in the linear predictors for both likelihoods.

Non-informative priors are a common choice when little expert knowledge is available. This approach is used here to define prior distributions over the fixed effects and the intercepts. Specifically, a zero-mean Gaussian prior with a precision of 0.001 is chosen for all fixed effects and intercepts. Prior information with different strengths can also be defined using the penalised complexity (PC) prior approach (Simpson *et al.*, 2017). This procedure penalises excessively complex models by placing an exponential prior on a distance to a simpler base model, which helps to stabilise the estimation. Priors then shrink model components toward their base models, thus preventing over-fitting. For the precision τ of the log-size observations, a weakly informative prior is set such that the probability of observing a standard deviation $(1/\tau)$ larger than the empirical standard deviation for the response is 0.01. For the precision parameters of the random walks of order 1, weak prior distributions were set where the probability that the standard deviations $(1/\tau_{0,k})$ corresponding to SU_A, SU_{MD} and TWI are greater than 5, 0.5 and 0.1, respectively is 0.01. For the precision parameters of the remaining random walks of order 1, relatively weak prior distributions were set where the probability that the standard deviations of the random walks are greater than the empirical standard deviation of the response is 0.01. Due to the lack of prior knowledge, this choice made sense since, a priori, the effects on the linear predictor can only be interpreted relative to the likelihood and the implicit scaling in the likelihood.

For both Matérn range parameters, the selection was guided using the empirical variogram and set a prior distribution where the probability that the range is smaller than 25km is 50%. Finally, for the Matérn marginal variances, a prior distribution where the probability that the variance is larger than 0.25 is 50% was set. As the triangulation mesh is vital when representing the spatial process, the sensitivity of the parameter estimates were studied in changes made to the mesh. It was found that a mesh with $4,600$ nodes was a good compromise between the accuracy and speed of the algorithm (see Figure 2.2c).

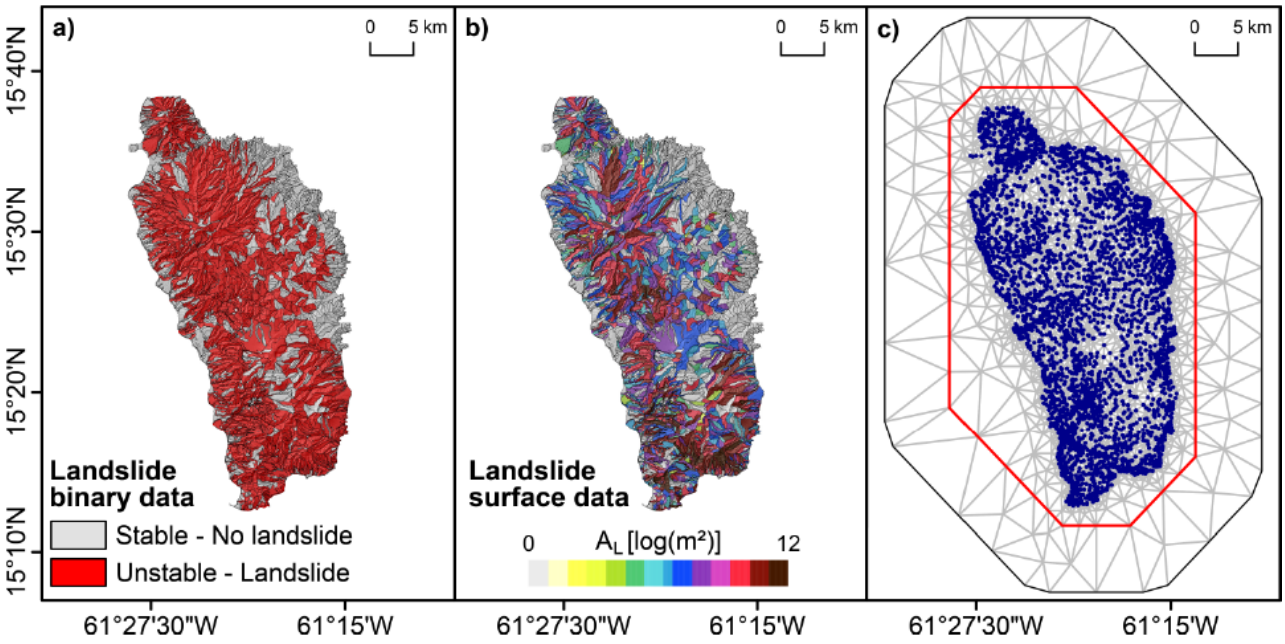

Figure 2.2: a) Observed landslide presence/absence data; b) Observed landslide planimetric area, aggregated at the SU level as the sum of all landslides; c) Triangulation mesh (gray) with SU centroids (blue) used to discretise Dominica Island and fit the spatial effect $u(\mathbf{s})$ in Equation 2.1. The inner boundary (red) delimits the island, whereas the extension to the outer boundary (black) avoids possible boundary effects.

2.4 Results

2.4.1 Model-level results and interpretation

This section presents the results of the hurdle model in terms of the statistical findings and the implications for landslide hazard assessment.

Figure 2.3 shows the posterior means and corresponding 95% credible intervals of the fixed effects

for the Bernoulli and Gaussian models. The selected covariates show relatively moderate positive and negative influences on landslide log-sizes and occurrences. The extent to which a covariate, significant or not, contributes to the model can be summarised by the size of the posterior mean regression coefficient. For the Bernoulli model, out of the 29 covariates used linearly, eight appeared to be significant; $SU_{MD/A}$, ENμ, PLCμ, PLCσ, PRCσ, RSPσ, SLOσ and IPL, with mean regression coefficients of -0.731, 0.228, -0.190, 0.693, -0.815, -0.217, 0.521 and 0.793, respectively. From here, the contribution becomes less distinguishable.

For the Gaussian model, out of 19 covariates used linearly, eight appeared to be significant, which indicates that the model is 95% certain of the role (either positive or negative) of the given covariate with respect to the landslide log-size. From the plot in Figure 2.3 it is shown that the significant covariates are $SU_{MD/A}$, $SU_{MD/\sqrt{A}}$, ELσ, PLCμ, PRCσ RSPμ, SLOσ and TWIσ, with mean regression coefficients of -0.835, 0.165, 0.146, -0.269, -0.244, -0.322, 0.439 and -0.148, respectively. It is important to notice that even if an effect is not significant, the effect size of the corresponding covariate may be large. Therefore, non-significant effects do not imply that the model is not influenced by the corresponding covariates (Lombardo *et al.*, 2021).

Figure 2.4 displays the posterior means and corresponding 95% credible bands of the random effects for both models. One can see the highly non-linear influence of most of these covariates on landslide log-sizes and occurrences. For example, looking at the plot for landslide susceptibility, it is clear that SU_A has a moderate non-linear effect, with a positive effect peaking at approximately $0.67km^2$. ELμ displays a slight concave effect, with a positive effect on landslide susceptibility between 0 and 0.4 m/1000. Also shown is that PRCμ has a relatively linear and mild effect over landslide occurrences, while SLOμ is among the most significant non-linear effects with very narrow credible intervals. As expected, steeper slopes are more at risk.

For landslide sizes, SU_A has a moderate non-linear effect, with a negative effect for smaller areas and an increasing and eventually positive effect for larger areas. SU_{MD} seems to have a mild concave and then convex effect on landslide log-size before becoming relatively linear. D2Sμ is relatively linear when considering the credible bands, with negative effects at small distances. ELμ has an overall negative effect on landslide size except for the range 0 - 0.4 m/1000. SLOμ steepness is again among the most significant effects with very narrow credible intervals. As expected, larger slopes - those above 30 degrees - are more at risk. Finally, TWIμ has a convex relation with landslide log-size.

Figure 2.5 show the posterior mean estimates at the mesh nodes of the spatial fields for both models. The spatial field covers almost the entirety of the island for the Bernoulli model, except for

the north-eastern coastlines. This is well in line with observations of Maria passing over Dominica and discharging such an amount of rainfall that most of the shallow material draping over the island has become prone to fail. The same cannot be observed in the case of the spatial field estimated for the Gaussian model. In this case, large landslides appear to share a high degree of spatial dependence in the centre of the island.

To assess the ability of the Bernoulli model to classify landslides correctly, the left panel of Figure 2.6 shows the receiver operating characteristic (ROC) curve. It is constructed by plotting the true positive rate (TPR, also called sensitivity) versus the false positive rate (FPR, also calculated as 1-specificity). The TPR boils down to the number of unstable SUs that have been correctly classified divided by the total number of unstable SUs. As for the FPR, this measure is calculated as the number of misclassified stable SUs divided by the total number of stable SUs. The input to the Bernoulli model has a dichotomous nature, whereas the output is a probability. Therefore, to match the output to the input, one needs to use a probability threshold. The way a ROC curve is constructed is by selecting a large number of these thresholds and storing the TPR and FPR at each iteration. The greater the area under the ROC curve (AUROC), the better the model and its classification abilities (Zou *et al.*, 2005). The Bernoulli model has an AUROC of 0.927, so it is safely affirmed that the model does an excellent job of distinguishing between both classes.

To assess the goodness of fit of the Gaussian model, the middle and right panels of Figure 2.6 show a histogram of the probability integral transform (PIT) values and a plot of observed versus predicted values, respectively. PITs are commonly used to assess model calibration, i.e., the statistical consistency between the predictive distribution and the observations (Gneiting *et al.*, 2007). If a model is well-calibrated, then the observations should be indistinguishable from a random draw from the model. For a large number of observations, the PITs histogram serves as a tool to empirically check for uniformity. As expected, large landslide sizes seem to be underestimated by the Gaussian model.

Nonetheless, the PITs in Figure 2.6 are not too far from the way a histogram of uniform numbers might look, and it seems fair to assume that the model is decently calibrated. The observed versus predicted values plot shows a relatively good performance for moderate landslide sizes, although there seems to be a greater bias for landslide sizes below the first quartile.

The advantage of a statistical model over other data-driven approaches is that the association between dependent and independent variables can be clearly interpreted. This is particularly useful for examining the geomorphological reasonability of a given model. Below, a few examples are presented of the interpretation for those covariates that behave close to the assumptions from the mechanical

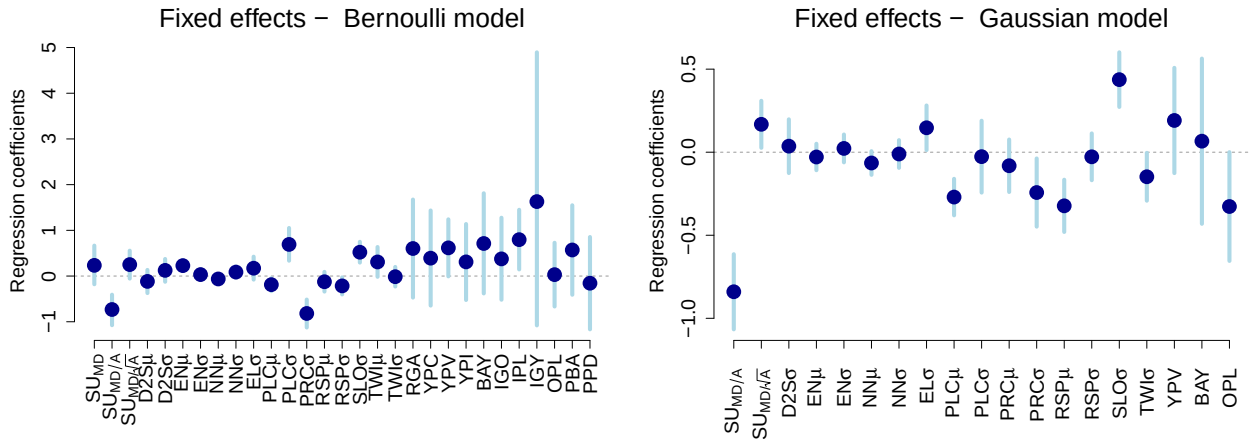

Figure 2.3: Posterior means (dots) of fixed linear effects (except the intercept) with 95% credible intervals (vertical segments) for the Bernoulli and Gaussian models. The horizontal grey dashed lines indicate no contribution to the landslide occurrence or size, respectively.

perspective. In this case, the interpretation should cover a dual aspect of the hurdle model, both from the landslide occurrence and size perspectives. As previously mentioned, Figure 2.3 reports the posterior distribution of the regression coefficients estimated for the Bernoulli and the Gaussian models. There, among the lithological classes, ignimbrites (IGO, IPL, IGY) positively contribute to increasing the mean probability of landslide occurrence, with only IPL being significant among the three. From an interpretative standpoint, ignimbrite is a pyroclastic flow deposit largely comprised of pumice with subordinate ashes. Thus, a positive contribution to the landslide occurrence is reasonable because its mineral structure is prone to weathering, and it is well known to promote slope instabilities (e.g., Chigira and Yokoyama, 2005). An analogous situation can be seen in gravel and alluvium materials (RGA). These are unconsolidated deposits, inherently susceptible to slope instabilities anytime the landscape evolution has set them to drape over steep topographies.

Interestingly enough, most of the lithotypes selected for the Bernoulli model do not appear in the Gaussian case. There, Older Pleistocene volcanics are also reported with a significant and negative contribution to the estimated landslide sizes. Irrespective of the Bernoulli or Gaussian framework, no land-use class plays a role in explaining the landslide occurrence or size.

Another point in common between the Bernoulli and the Gaussian models is the role of $SU_{MD/A}$. This covariate expresses how elongated a given SU is. The larger the value, the more stretched the SU appears, whereas the smaller the value, the more rounded the SU is. Thus, the posterior mean negative value reported for both models may indicate that narrow SUs are not only less prone to fail,

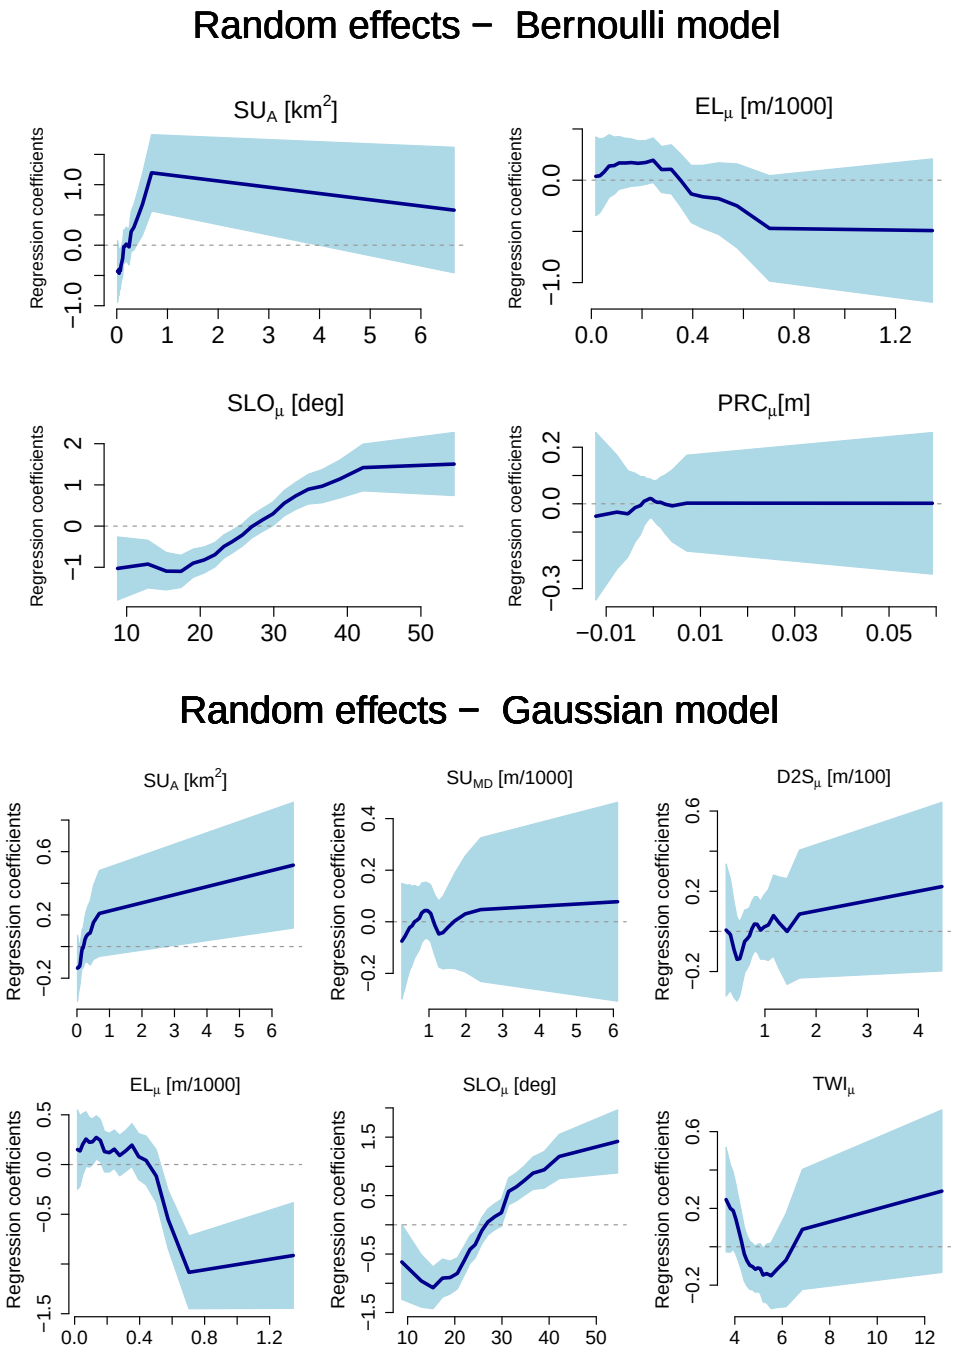

Figure 2.4: Posterior means (dark blue curves) of random effects with 95% pointwise credible intervals (light blue shaded area) for the Bernoulli and Gaussian models.

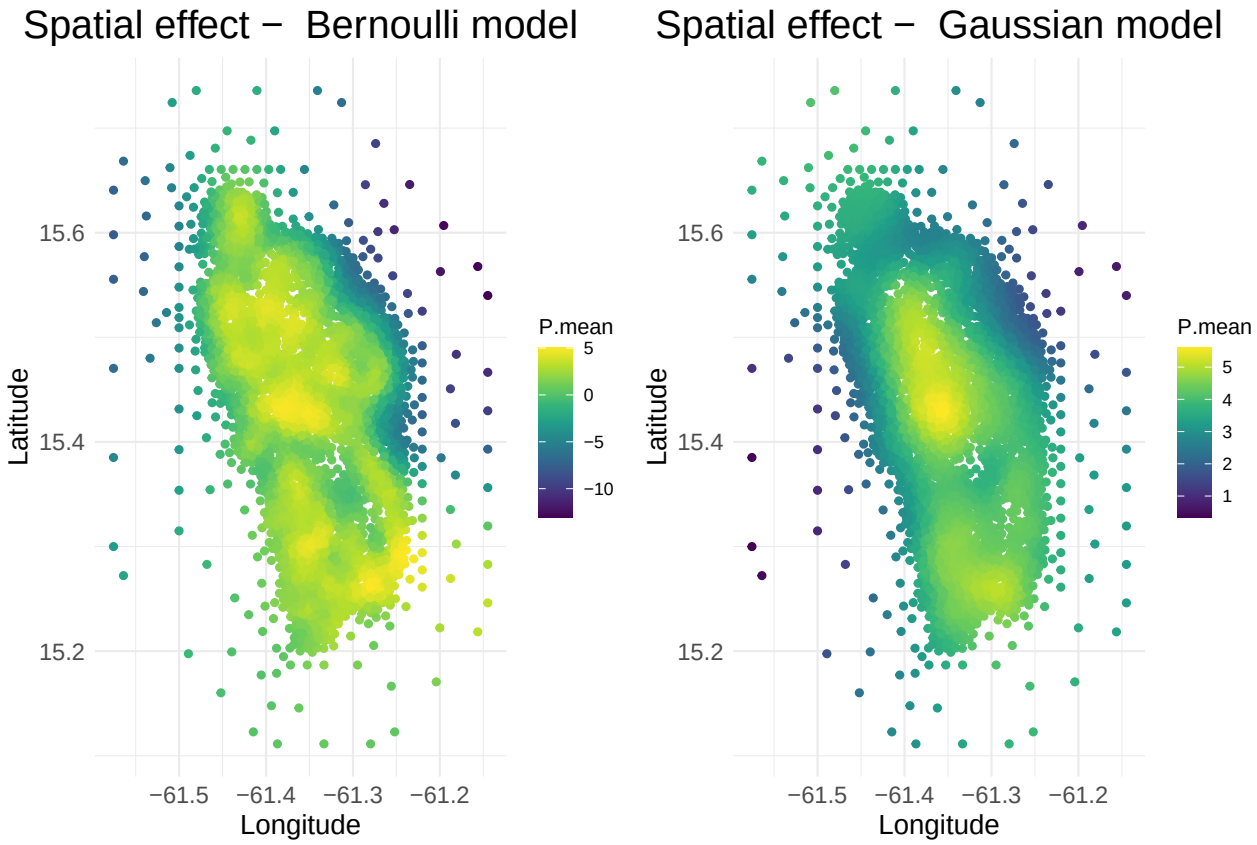

Figure 2.5: Posterior means of the spatial field for both models.

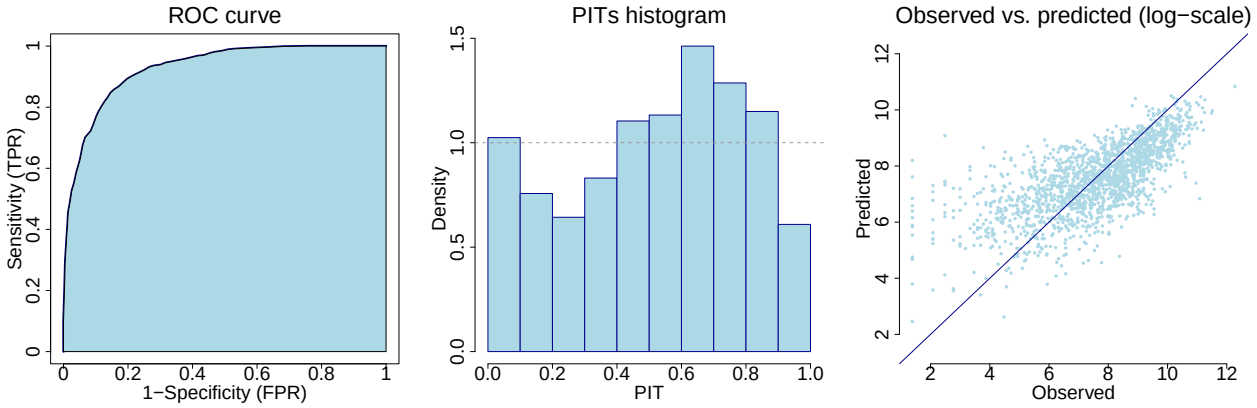

Figure 2.6: ROC curve for the Bernoulli model, PIT histogram, and observed versus fitted plot (in log-scale) for the Gaussian model.

but also the lesser availability of material does not allow for large landslides to be generated.

As for the non-linear effects shown in Figure 2.4, observed are two noteworthy behaviours. SU_A appears to influence both the Bernoulli and Gaussian models with a negative effect for very small SUs, while the contribution becomes increasingly positive for larger SUs. Similarly to the elongation/roundness ($SU_{MD/A}$) effect mentioned above, this can be interpreted as if larger SUs can host more and larger landslides compared to very small SUs. There, the conditions required for a failing mechanism to occur are much more unlikely to manifest simply because of the smaller extent. The other interesting effect corresponds to SLOμ. In both models, the slope steepness behaves as a sigmoidal function. This is likely because no shallow flow-like landslide can occur at low steepness values. Conversely, at medium steepness values, there is a sudden increase in the effect of this covariate, which is to be expected, up to the point where an asymptotic level is reached or even a decrease. A decrease can be interpreted as if very steep topographic conditions cannot host soil, which is washed away by normal erosional processes. Thus, the absence or near-absence of soil implies that no landslide can manifest or that, at best, a small one will mobilise the thin detrital layer draping over the stable bedrock.

Figure 2.7 presents the findings in terms of posterior means and width of the 95% credible intervals (CI) for landslide susceptibility and sizes (in log-scale) across the island. There, large portions of the areas characterised by higher landslide probability (central north and southern coastlines) coincide with areas where larger landslide sizes are expected. Conversely, the north-eastern coast seems consistently small in landslide size and probability of occurrence.

From the uncertainty plots, areas with higher occurrence probability (0.86 to 1) have overall small 95 % CI width, implying that these areas are well estimated with relatively small uncertainty. It is also shown that areas with moderate occurrence probability (0.47 to 0.86) display higher uncertainty, while areas with the smallest occurrence probability (0 to 0.47) show a mixture of small and moderate uncertainty. Large landslide sizes are estimated with relatively small uncertainty, and the higher uncertainties seem to occur where the landslides are small to moderately sized.

2.4.2 Unified landslide hazard assessment

The Bernoulli model addresses the island's susceptibility under extreme conditions, such as those induced by Maria. It estimates how prone a given location is to host slope failures, but it is blind to how large landslides may become once they initiate and propagate downhill. To compensate for this limitation, the Gaussian model provides information on the expected size of landslides per slope unit,

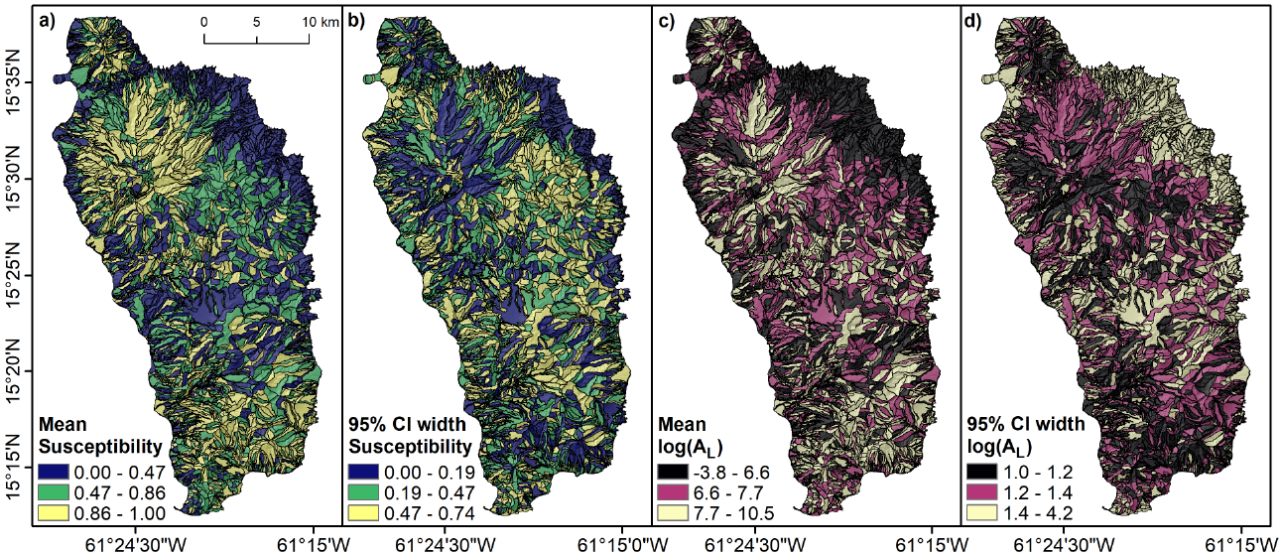

Figure 2.7: Map of Dominica showing posterior mean and width of the 95% credible interval for landslide susceptibility (panels (a) and (b)) and landslide log-size (panels (c) and (d)).

again under extreme meteorological stress. But once more, the model is blind to which slope was effectively prone to fail. So far, the results of the two models have been implemented and presented independently of each other. This framework is already more than enough to satisfy the requirements of the most accepted landslide hazard definition (see Guzzetti *et al.*, 1999), taking aside the temporal dimension. Nevertheless, much more can be done to provide for the first time a unified version of the landslide hazard for the purely spatial context. In this subsection, a new and unified (data-driven-specific) landslide hazard assessment is described by combining the two elements of the proposed hurdle model.

The motivation to provide a unified framework stemmed from the need to provide end-users with spatially distributed information regarding how likely each slope would be to release specific landslide sizes if expected to be unstable under analogous extreme conditions to those brought by Maria. Small landslides should carry the least hazard; thus, they should be of little interest. Conversely, as the landslide size increases, the expected hazard should proportionally follow. So, to quantify the hazard for moderate and relatively large landslides within the range of observed landslide sizes, the two parts of the hurdle model are used to compute the probability of observing landslide sizes above the empirical 50%, 75%, 90% and 95% quantiles. Specifically:

$$\Pr(\log\{\mathrm{A_L}(\mathbf{s})\} > u) = \Pr(\log\{\mathrm{A_L}(\mathbf{s})\} > u \mid \mathrm{O_L}(\mathbf{s}) = 1)\Pr(\mathrm{O_L}(\mathbf{s}) = 1), \tag{2.2}$$

for $u = \hat{F}^{-1}_{\log(\mathrm{A_L})}(0.5), \hat{F}^{-1}_{\log(\mathrm{A_L})}(0.75), \hat{F}^{-1}_{\log(\mathrm{A_L})}(0.90)$ and $\hat{F}^{-1}_{\log(\mathrm{A_L})}(0.95)$, where $\hat{F}_{\log(\mathrm{A_L})}$ is the empirical cumulative distribution function of landslide log-sizes. Note that this procedure can be performed for any landslide size u of interest. The focus here is on the empirical 50%, 75%, 90% and 95% quantiles to illustrate how the model can be used to predict the probability of exceeding medium to large landslides. The exceedance probabilities $\Pr(\log\{\mathrm{A_L}(\mathbf{s})\} > u)$ in (2.2) and their uncertainties can be easily computed using posterior samples from both the Gaussian and Bernoulli models. Specifically, a Monte Carlo (MC) procedure is followed and $N = 500$ samples are generated from the posterior predictive distributions (PPDs) of the Gaussian and Bernoulli models. The PPDs account for uncertainty in the data and the model fitting. The empirical estimates are then computed of the probabilities on the right side of (2.2) to obtain one estimate of $\Pr(\log\{\mathrm{A_L}(\mathbf{s})\} > u)$. Finally, this procedure is replicated $M = 1,000$ times to obtain MC mean estimates and confidence intervals for $\Pr(\log\{\mathrm{A_L}(\mathbf{s})\} > u)$, computed as the 2.5% and 97.5% quantiles of the MC estimates. Figure 2.8 shows plots of the exceedance probabilities and the width of the associated 95% confidence intervals for the four quantiles detailed above. It is interesting to observe how the exceedance probability of the landslide sizes changes from one hazard map to the other. As one should expect, the slopes prone to releasing median landslide sizes are quite numerous, and their number decreases towards larger landslides. To briefly touch on risk perspectives (although not explicitly integrated into the hurdle model), landslides greater than 90% of the landslide area distribution are particularly likely to occur in the southernmost sector of the island. There, the village of Berekua appears to be potentially vulnerable, the location indicated approximately on Figure 2.8 by the yellow star. The same situation can be seen slightly northwestward for the much larger settlement of Roseau, the location indicated approximately on Figure 2.8 by the white star. This type of consideration would not be possible in the simple binary case, where the corresponding susceptibility map highlights most of the island as unstable (see Figure 2.7a). Note that the confidence intervals in Figure 2.8 measure the accuracy of the Monte Carlo approximation and should not be interpreted as a measure of the dispersion in the posterior predictive distribution of the exceedance probabilities.

2.5 Discussion

A complete landslide hazard assessment should address three components: where, when (or how frequently) and how large landslides may be. The three components mentioned above have always been addressed separately in the scientific literature produced so far (for non-physically-based models).

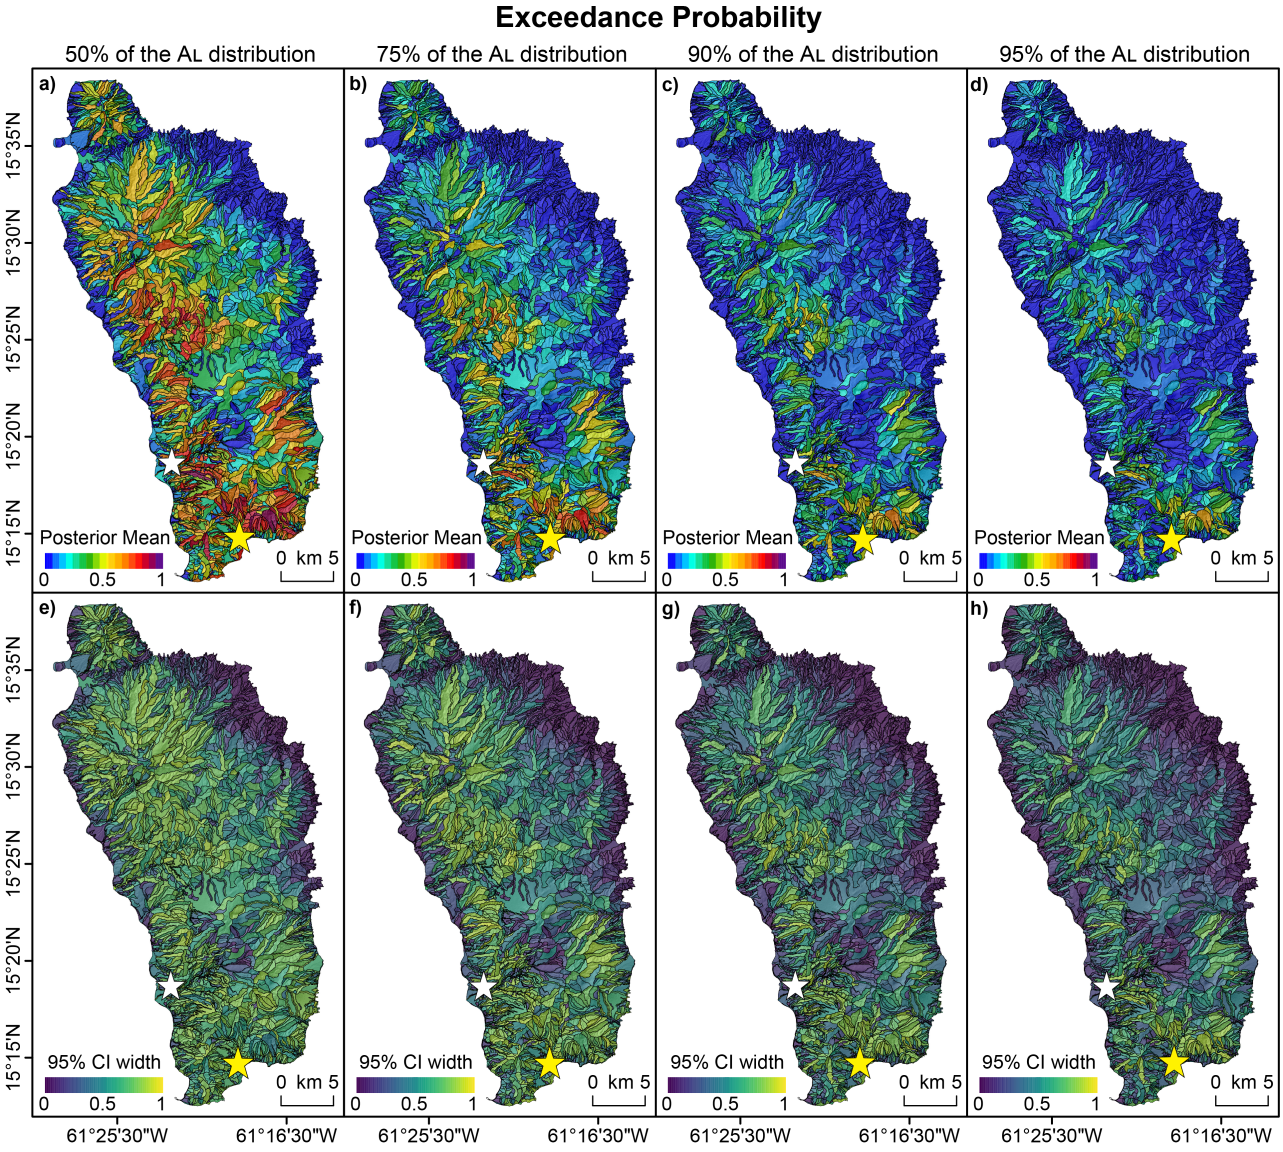

Figure 2.8: Maps displaying exceedance probability estimates (top) and a measure of their uncertainty (bottom) for landslide sizes above the 50%, 75%, 90% and 95% empirical quantiles. The computations were based on $N \times M = 5 \times 10^5$ predictive posterior samples of the Gaussian and Bernoulli models. The white and yellow stars indicate the settlements of Roseau and Berekua, respectively.

In this work, two of the three were combined, leaving the temporal characteristic out of the scope because of the limited access to long-time series of landslide inventories triggered in response to hurricanes within the Island of Dominica. Despite this limitation, the model proposed represents a substantial improvement with respect to traditional presence/absence susceptibility models because those are blind to the actual threat that a landslide may pose, depending on its size.

This could be a stepping stone for further experimentation in the hope that one day this unified hazard line of research may impact international guidelines for disaster risk reduction. In this work, focus was on the landslide inventory generated by Maria because of the large sample size and the availability of data on a wide variety of geographical and geological variables, with no missing values. However, further effort can be made to implement an extension in space-time of the hurdle framework.

Another potential improvement would be to extend this hurdle model to different modes of slope failures and propagation. Currently, only shallow flow-like landslides are modelled that either started as debris slides or flows, which generally evolved into debris flows due to the high water content of the moving mass. However, extensions of the model could be implemented to distinguish various landslide classes, including deep-seated ones, which have completely different failure mechanisms and propagation behaviour.

Ultimately, the integration of actual rainfall data is also envisioned. Unfortunately, in the specific case of Maria, no reliable rainfall data is available. However, other inventories in data-rich conditions could help integrate rainfall into the modelling strategy, which in turn could enable testing for a near-real-time application of the hurdle model.

Chapter 3

A data-driven landslide susceptibility model for Scotland

This chapter presents the first data-driven model for landslide susceptibility in Scotland, developed using a national landslide inventory compiled by the British Geological Survey. We begin by fitting a Bernoulli model to estimate the probability of landslide occurrence, providing a comparison to an existing heuristic model. Building on this, we propose an extension using a log-Gaussian Cox process model. The work in this chapter has been published as

Bryce, E., Castro-Camilo, D., Dashwood, C., Tanyas, H., Ciurean, R., Novellino, A. and Lombardo, L., 2025. An updated landslide susceptibility model and a log-Gaussian Cox process extension for Scotland. *Landslides*, 22(2), pp.517-535.

3.1 Introduction

The British Geological Survey (BGS) are a renowned organisation for knowledge in geological conditions across the UK and I was fortunate to work with them on a project to mitigate landslide risk across Scotland. The project was funded through the National Centre for Resilience Project Grant Award (NCRR2023-001) and the BGS International NC program 'Geoscience to tackle Global Environmental Challenges', NERC reference NE/X006255/1 with the aims to update the National Landslide Database (NLD), build a new landslide prediction model, and implement this to produce exposure risk maps. My role involved building a new landslide prediction model, as the previous (named GeoSure) was built for assessment purposes based mainly on ground movement and subsidence across the UK due to different geological conditions. The differences in data availability

and modelling approaches between GeoSure and my proposed model are compared throughout this chapter. Perhaps the most notable difference is that the GeoSure modelled the causative factors of landsliding using a heuristic approach requiring expert judgement.

GeoSure is a susceptibility model with five classifications ranging from low to high susceptibility, but with no quantification assigned to these classes. I firstly proposed a landslide prediction model with a similar methodological framework to the case study of Dominica (see Chapter 2), which involves modelling the probability of observing a landslide, in a given geographical unit previously defined, by using a Bernoulli distribution and inference techniques described in Section 1.4.

Although an efficient way to model landslides, assigning a binary response (presence/absence) at the given geographical unit removes the count aspect of landslides, as there could be multiple landslides in the unit, and we are only using a binary classifier. Therefore, a log-Gaussian Cox Process, Section 1.6, was proposed in order to account for this. The two model results are compared in Section 3.4, and interpretability and usefulness are examined for cross-validation measures in Section 3.5.

3.2 Data

The target of our study is Scotland, a main landmass surrounded by multiple smaller islands covering an area of $\sim 80,000\text{km}^2$ (Figure 3.1). Most of the British landmass is dominated by a gentle topography with a slope angle of less than five degrees, characterising 90% of the total territory (Cigna *et al.*, 2014). Conversely, the Scottish landscape is characterised by over 50% upland environments formed as a result of an interaction between glacial incisions and post-glacial isostatic uplift (Firth and Stewart, 2000). Scotland's diverse bedrock formations are covered by a thin or patchy cover of superficial deposits such as glacial till, hummocky morainic deposits as well as weathered bedrock. Notably, most shallow flow-like landslides take place in such superficial materials (Palamakumbura *et al.*, 2021). The impact of ice erosion has created several recognisably distinct landscapes across Scotland including the western ice-scoured landscape, weathered bedrock and solifluction deposits in the far east, ice-scoured lowlands and extensively modified valleys, troughs, and mountains. Land-use in Scotland is quite homogeneous, with a large percentage of its surface dedicated to agriculture (∼70%), with woodland corresponding to most of the remaining cover, according to the CORINE Landcover map (CLC) As for the built-up areas, these occupy a minor extent, with urban classes accounting for less than 3%. The main cities are spread across the central belt, and communications

among centres are ensured by a road network relying on a few main transport arteries. In places, the lack of diversions away from these key routes makes them sensitive to the occurrence of landslides. The Scottish road and rail networks are regularly affected by debris flows, with the most well-known example of this being the Rest and Be Thankful (RABT). The RABT was closed in October 2023 after being hit by seven landslides in just a few days, whilst in the same storm event, ten people were airlifted to safety after becoming stuck between landslides on the A83 and A815. Previous heavy rainfall events have led to a debris flow that affected the Fort William to Mallaig train line as well as the A830 highway in 2016 in an area previously considered to be of low debris flow susceptibility (Palamakumbura *et al.*, 2021). Whilst each one of these events was relatively small in scale, the impacts on local communities and businesses can be significant. The village of Ardfern, for instance, remained inaccessible a month after a 6000-tonne landslide blocked the A816, rendering it impassable. The potential for these types of landslides to increase with future climate change highlights the need for modern, up-to-date LSMs.

The BGS National Landslide Database (NLD) is a continuously updated source of landslide information. For this study, we extracted Debris Flow (DF) locations because they are the most common landslide type in Scotland, and they cause the largest impact on infrastructure routes. The NLD has changed the way landslides have been collected. Initially, BGS would collect information sourced from news reports, individual and transport institution reports, whilst recent developments have seen the use of social media and earth observations from satellite scenes (Pennington *et al.*, 2015). This combined search is meant to ensure that any potential bias in the spatial distribution of landslides is minimised. Understanding bias is important when dealing with a spatially distributed process, such as landslides. The inclusion of earth observations is intended to limit the skew towards transportation routes and urban areas that can be produced by collecting data purely from social media posts and reports from transport infrastructure operators. To further understand the implications of such potentially biased sources, we suggest reading the work of Lima *et al.* (2021) or Lin *et al.* (2021). To avoid propagating such biases, one could use two potential solutions. The first is in introducing bias-related covariates in the model fitting, which are then zeroed out in the model prediction phase (Lima *et al.*, 2017). The second and alternative solution is to solve the issue at the source by introducing an independent mapping procedure. The strategy of the BGS includes the use of freely available Sentinel-2 satellite images and an approach similar to NASA's Sudden Landslide Identification Product (SLIP) tool (Fayne *et al.*, 2019) to automatically map potential debris flow locations. This combined approach ensures that the NLD reflects the standards of quality (Galli *et al.*, 2008)

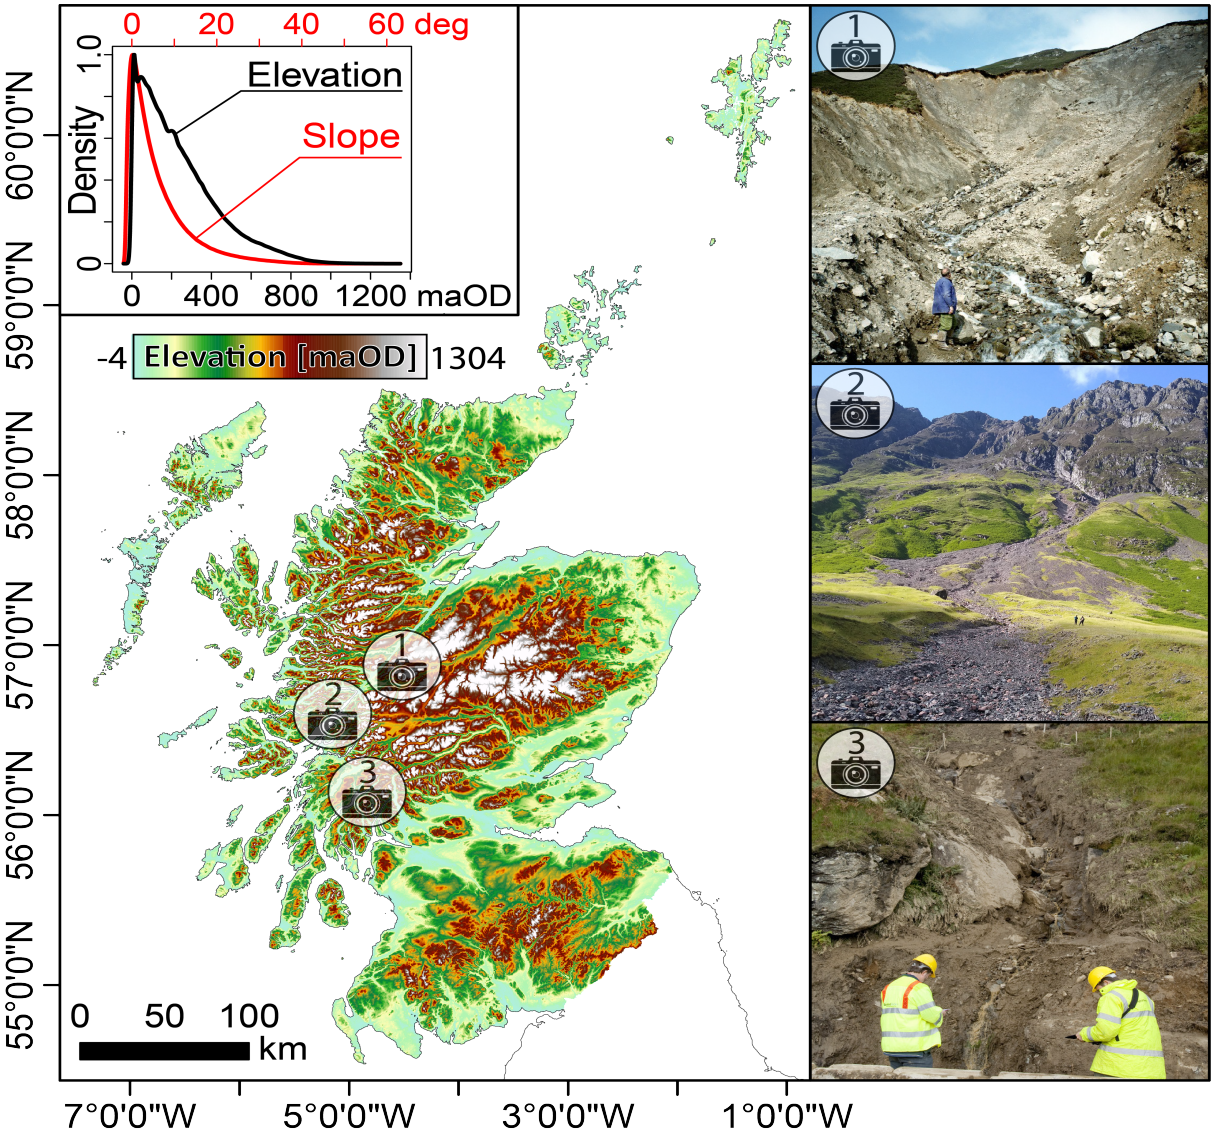

Figure 3.1: Terrain overview of Scotland summarised both cartographically and as the two probability density plots for elevation and steepness. The three photos were taken from the BGS field collection of debris flows. NEXTMap Britain elevation data from Intermap Technologies; Photo number 1 is a BGS image P001177© UKRI 1990; Photo number 2 is a previous BGS image; Photo number 3 is a BGS image P757938© UKRI 2009.

and completeness (Tanyaş *et al.*, 2022) required for a suitable landslide prediction modelling protocol. The resulting inventory features 1,854 DFs across the Scottish mountainous terrain. These DFs have been digitally recorded with a point whose X and Y coordinates correspond to the highest position visible on the source area. The reason behind such a choice is to represent the most likely location where the failure was initiated (Scheip and Wegmann, 2022). Notably, this may be an approximation because laboratory experiments have proven that DFs may exhibit retrogressive behaviour (Sosio *et al.*, 2007). However, as one faces the limitation of being only able to examine the scarp left by the DF, choosing the highest point along a polygon perimeter is the most reasonable approach (Lombardo *et al.*, 2014). Nevertheless, this level of detail is to be accounted for whenever the mapping unit of interest is expressed at high resolution. At the level of a coarse spatial partition, the assignment of a stable/unstable label would not change. In this research, we opted to partition Scotland into Slope Units (SUs) and therefore, no substantial changes are expected for both the susceptibility and intensity models.

To represent the DF information over space, we chose the SU partition of Scotland, which can be seen in Figure 3.2. A Slope Unit (SU) encompasses the geographic space between streams and ridges (Amato *et al.*, 2019), and a number of analytical tools to extract them from digital elevation models (DEMs) have been developed over time. From the first inception by (Carrara *et al.*, 1991), SUs have been manually mapped (Guzzetti *et al.*, 1999) and later obtained via the Inverse DEM method (Turel and Frost, 2011). Recently, though, a robust computational scheme has been introduced by (Alvioli *et al.*, 2016) in the form of a GRASS GIS (Neteler and Mitasova, 2002) script named `r.slopeunits`. In this work, we opted to use `r.slopeunits`, parameterising it with the following values:

- Flow accumulation threshold = 1,000,000
- Circular variance = 0.3
- Minimum area = 25,000
- Clean size = 10,000

The resulting procedure produced a total of 153,282 SUs, whose geographic overview and frequency area distribution are shown in Figure 3.2.

Notably, SUs are irregular polygonal objects whose spatial extent largely exceeds the resolution of common terrain and thematic covariates. For this reason, SUs require an aggregation step to

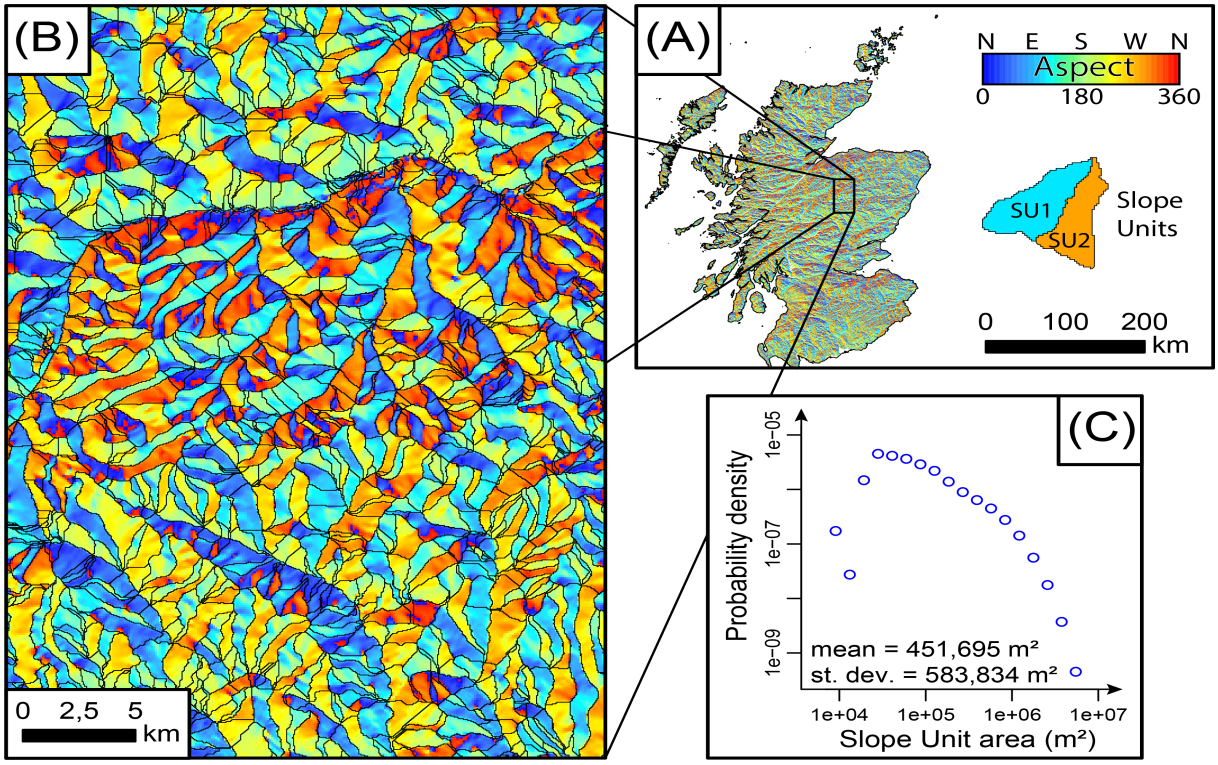

Figure 3.2: Panel A shows the aspect distribution over the entirety of the Scottish landscape. Panel B highlights the Slope Unit delineation over a small area to provide an overview of the partition and the typical mapping unit sizes. Panel C is the frequency-area distribution of the Slope Units over the whole of Scotland.

express both dependent and independent variables for each polygon. The dependent one corresponds in our case to the landslide presence and absence, to be assigned at the intersection between landslide identification points and SUs. This aspect concerns the susceptibility model only. In fact, for the LGCP model, the point pattern theoretical foundation requires the locations to be kept with their original details (Bagchi and Illian, 2015). The aggregation at the SU level becomes part of the modelling protocol only as a post-processing routine. To assign covariate values, we compute the mean and standard deviation of each variable within the boundaries of an SU polygon. These summary statistics are denoted by the suffixes μ and σ in Table 3.1, which lists all covariates selected to explain the spatial distribution of landslides in Scotland. An exception is made for rainfall, for which no summary statistics were calculated due to the coarse spatial resolution (1 km) of the precipitation data. For the lithological covariates, aggregation was carried out by identifying the predominant class within each SU.

Table 3.1: Summary of the BGS dataset, including the responses and initial covariates' set.

Variable	Acronym	Type	Units	Resolution
DF inventory occurrence	$O_{DF}(\mathbf{s})$	Binary response	0=No, 1=Yes	N/A
DF inventory count	$Count_{DF}(\mathbf{s})$	Continuous response	Unit-less	N/A
SU area	SU_A	Continuous explanatory	m^2	50m and 5m
Local relief, 1000m buffer	LR	Continuous explanatory	m	10m
Annual precipitation mean (1999-2019)	$Prec_\mu$	Continuous explanatory	mm	1000m
Annual precipitation maximum (1999-2019)	$Prec_{max}$	Continuous explanatory	mm	1000m
Planform curvature mean and standard deviation	PLC μ and σ	Continuous explanatory	1/m	10m
Profile curvature mean and standard deviation	PRC μ and σ	Continuous explanatory	1/m	10m
Slope mean and standard deviation	SLO μ and σ	Continuous explanatory	Degrees	10m
Quaternary domains	*Quat*	Categorical explanatory	Unit-less	1:10,000-1:1,000,000
Superficial deposits	*Super*	Categorical explanatory	Unit-less	1:625,000
Bedrock	*Bedrock*	Categorical explanatory	Unit-less	1:625,000

3.3 The models

As mentioned above, two modelling archetypes were implemented: to model landslide susceptibility as per international standards, the probability of observing at least one DF in a SU was predicted by using a Bernoulli distribution. For the LGCP, the DF rate of occurrence is modelled per SU by using a Poisson distribution with a random intensity function that approximates the LGCP likelihood of the landslide points distributed across the space. In both cases, it is assumed that the observations (presence/absence in the susceptibility case or counts in the LGCP case) are conditionally independent given a latent Gaussian process (see Section 1.3), where these models can flexibly capture local correlation structure and uncertainty. As a result, the covariates can be modelled flexibly in terms of their influence on DF occurrence or intensity.

The LGCP is a spatial Poisson process with a spatially varying intensity, $\lambda(\mathbf{s})$, and modelled as a Gaussian process in the log scale (Bachl *et al.*, 2019). The intensity of the point process is linked to

the linear predictor as follows:

$$\log \lambda(\mathbf{s}) = \eta(\mathbf{s})$$
$$\eta(\mathbf{s}) = \alpha + \sum_{j=1}^{J} \beta_j x_j(\mathbf{s}) + \sum_{k=1}^{K} f_k(z_k(\mathbf{s})) + u(\mathbf{s}),$$

where $\mathbf{s} = \{s_1, ..., s_n\}$ are the spatial units that discretise the study region, α is an intercept and $(x_1(\mathbf{s}), ..., x_J(\mathbf{s}))^T$ are environmental covariates that might influence, in a linear way, the variable of interest, with coefficients $\boldsymbol{\beta} = (\beta_1, ..., \beta_J)^T$. The functions $f = \{f_k(.), ..., f_K(.)\}$ are non-linear effects defined in terms of environmental covariates $(z_1(\mathbf{s}), ..., z_K(\mathbf{s}))^T$. The specific form of the functions $f_k(.)$ can be those of a Gaussian random walk of order 1 or 2, for example (RW1; Krainski *et al.* (2018)). Lastly, the term $u(\mathbf{s})$ is a Gaussian Random Field (GRF) approximated via the stochastic partial differential equation approach (SPDE; Lindgren *et al.*, 2011).

We model DF susceptibility with a Bernoulli distribution, thus $y(\mathbf{s}) = \mathrm{O}_{\mathrm{DF}(\mathbf{s})} \in \{0, 1\}$ and $y(\mathbf{s}) \mid \eta_{\mathrm{Bern}}(\mathbf{s}) \equiv \mathrm{Bern}(\mathrm{p}(\mathbf{s}))$, where $\mathrm{p}(\mathbf{s}) = \Pr\{\mathrm{O}_{\mathrm{DF}}(\mathbf{s}) = 1\}$. The probability $\mathrm{p}(\mathbf{s})$ is related to the linear predictor $\eta(\mathbf{s})$ through the logit link, so

$$\mathrm{p}(\mathbf{s}) = \frac{\exp\{\eta_{\mathrm{Bern}}(\mathbf{s})\}}{1 + \exp\{\eta_{\mathrm{Bern}}(\mathbf{s})\}}. \tag{3.1}$$

The sum of the susceptibility model components, $\eta_{\mathrm{Bern}}(\mathbf{s})$, follows the general model structure of Equation 1.2, with its specific form depending on the influence of the covariates on the DF susceptibility. To find the most appropriate way to express the influence of each covariate, as well as whether each given covariate provides useful information to the model, we conduct a dual-stepped variable selection. This is performed by testing each covariate in a linear and nonlinear form, as well as introducing these two realisations as part of a standard stepwise forward procedure (Steyerberg *et al.*, 1999). This procedure calculates the DIC and WAIC. Out of all the covariates listed in Table 3.1, the selected ones and their specific form as part of the susceptibility model are detailed below in Table 3.2.

Table 3.2: Selected covariates to be used in the susceptibility model.

Fixed effects	**Random effects**	**Categorical effects**
SLO σ, Prec_{max}	LR, SLO μ	*Quat*

The spatial rate of DF occurrences per SU (DF intensity) is modelled via an LGCP. This model has a doubly stochastic nature consisting of an inhomogeneous Poisson point process whose random

intensity surface is expressed in the logarithmic scale, allowing it to be modelled with a Gaussian likelihood (Illian *et al.*, 2008). As outlined in Section 1.6, Illian *et al.* (2012) proposed a method to approximate the number of points in a grid cell and by assuming that the lattice of grid cells is fine enough, this converges to the true LGCP likelihood.

With this polygonal data, rather than defining a regular lattice over the domain and losing interpretability (recall the definition of the SU is to discretise the space in such a way that the conditions for landsliding make geological sense - we would like to maintain this interpretability), we use the SU partition as a proxy for the grid cell and say that the SU map of polygons is fine enough. This can be further broken down as the number of points in each SU following a Poisson distribution with its mean represented by the intensity of the cell. This intensity is then approximately equal to the area of each cell multiplied by the exponential value of the latent field in each cell.

This process ensures that the number of DFs occurring in a SU can be considered rather than simplified into the binary classification typical of susceptibility studies. In other words, the susceptibility case keeps the zeroes exactly in the same form as the LGCP. However, the positive value is compressed to one, denoting slope instability. Conversely, the LGCP framework allows modelling the numerosity of the slope failures rather than being limited to the presence/absence situation. The selected covariates and their specific form of entry underwent the same variable selection procedure described for the susceptibility model. Out of all the covariates listed in Table 3.1, the selected ones and their specific form as part of the LGCP model are detailed below in Table 3.3.

Table 3.3: Selected covariates to be used in the LGCP model.

Fixed effects	**Random effects**
SLO σ, PRCσ	LR, SLO μ, Prec_{max}

3.4 Results

This section reports the outcome of the modelling protocol. Due to the dual set of experiments, the susceptibility and intensity results will be presented separately, first by showcasing the covariate effects, then converting the model estimates into map form.

3.4.1 Susceptibility

The marginal effects of the covariates selected and displayed in Table 3.2 are presented in Figure 3.3. There, the influence of the local relief (LR) appears as the most dominant covariate among all the selected ones. LR is computed as the difference between the maximum and minimum elevation values within a single SU. Therefore, this is commonly interpreted as a proxy for gravitational potential, a property intrinsically linked to DF predisposition, see Iverson (1997). Specifically, the inferred pattern indicates a negative contribution of the relief up to 200m. From this elevation difference to approximately 250m, the LR contribution to the DF occurrence probability rapidly increases and transitions to an approximate linear and positive trend up to 600m.

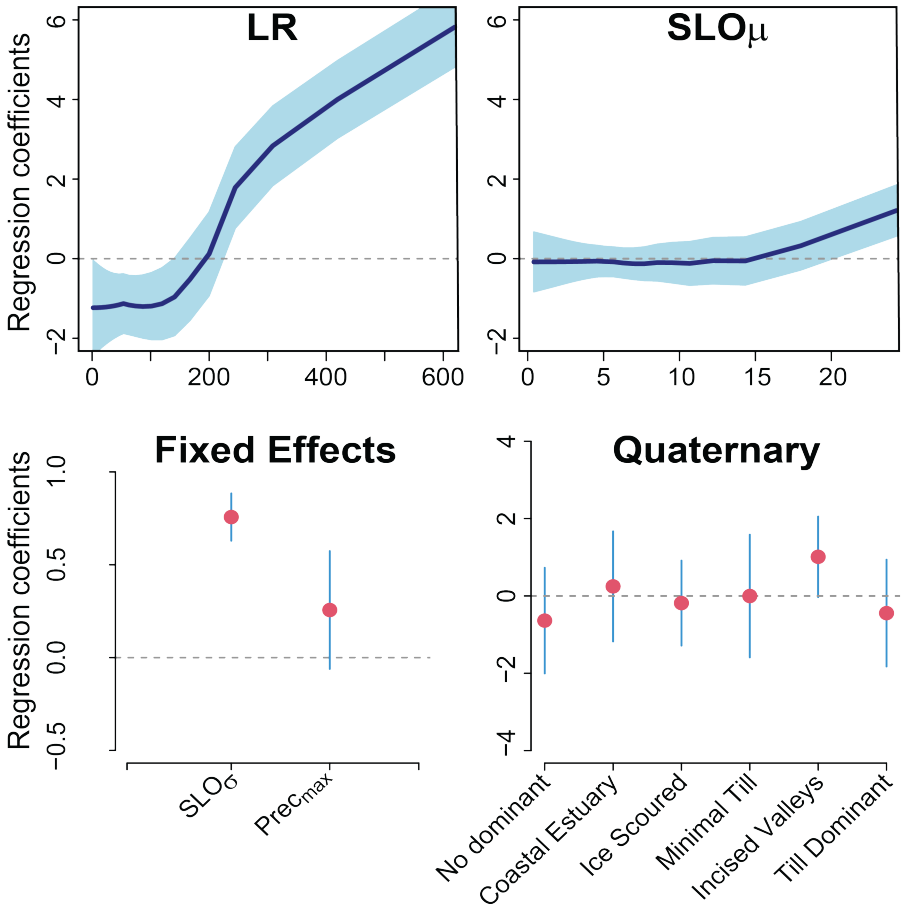

Figure 3.3: Posterior means (dark blue curves) of random effects with 95% pointwise credible intervals (light blue shaded area) (top row). Posterior means (dots) of fixed linear effects (except the intercept) with 95% credible intervals (vertical segments) and of categorical quaternary effects (bottom row). The horizontal grey dashed lines indicate no contribution to the DF occurrence.

Figure 3.3 additionally highlights the contribution of the average steepness per SU SLO μ. Slope steepness is also a measure of variation in elevation. The LR estimates differences over a large neighbourhood, whereas steepness values are computed as the first-order derivative between two adjacent grid cells. Therefore, the LR and SLO μ can be considered as the two sides of the same morphometric coin. A closer look at the latter marks a contribution to the susceptibility with a negligible effect up to approximately 15 degrees of mean steepness per slope unit. After which, the trend becomes positive and approximately linear up to the limit of 25 degrees. Interestingly, 21 degrees are empirically referred to by Iverson (1997) as a potential threshold for a slope to become prone to DFs. Moving to the fixed effects, two were selected as such. The first is the SLO σ, another parameter capable of capturing topographic roughness. Its contribution to the susceptibility model appears positive (mean $\beta_{\text{SLO}\mu} = 0.710$) and significant (95% of the regression coefficient distribution shares the same sign). A lower and still linear contribution is also estimated for the 20-year maximum rainfall amount computed per SU (mean $\beta_{\text{Prec}_{max}} = 0.009$). In this case, though, the covariate misses significance by a slight margin, with the 97.5 percentile of the regression coefficient distribution being markedly positive while the 2.5% appears negative. Nevertheless, the mean is still quite far away from the zero line shown in the plot, thus implying a non-negligible contribution to the model, on average, which is as expected for a covariate that should be linked to the DF genetic process. As DF records were not accompanied by their temporal information in the inventory metadata, we could only opt for a general meteorological representation of the Scottish landscape, rather than a precise measure of the trigger pattern in space and time. Ultimately, the bottom-right panel of Figure 3.3, presents the categorical effect of the morphology left by the last Quaternary glaciation retreat (addressed as Quaternary). Among all the landforms, only the Incised Valleys have been estimated with a positive and significant contribution to the DF occurrence probability. This is geomorphologically sound, and a result commonly retrieved in other DF susceptibility studies tailored to flow-like landslides in Scottish terrains (Ballantyne, 1986). All the remaining quaternary classes appear not to be statistically significant, nor do their regression coefficient appear to be large enough to cause notable variations to the susceptibility pattern, on average.

The sum of all mean susceptibility model components, together with the global intercept and after the logit transformation, produces the estimated susceptibility map shown in the left panel of Figure 3.4. Furthermore, the variability estimated for each of the regression coefficients shown in Figure 3.3, leads to the uncertainty estimates mapped in the central panel of Figure 3.4. What stands out is that the model produces susceptibility patterns for which the central sector of Scotland appears

to be largely prone to DFs. Conversely, the southernmost, easternmost and northernmost districts generally present non-susceptible characteristics. However, each of these districts is associated with a different probabilistic pattern when we include uncertainty-oriented considerations. The northern and eastern districts show very low mean susceptibility values associated with very low uncertainty values. Therefore, this is a portion of the landscape that is largely secure from a risk assessment perspective. As for the southern case, low susceptibility values are generally accompanied by high uncertainty levels. In turn, this may indicate a potential danger and require further attention rather than considering this district safe. Making such a consideration is crucial, and it is also the reason why Bayesian statistics is so widely adopted across virtually any scientific field. However, for a science such as geomorphology, intrinsically connected to environmental policies, producing separate maps and commenting on their relative patterns is not ideal. This is the case because traditionally, decision-makers do not have formal statistical training and, at times, a geoscientific one (Betcherman 1993). Therefore, reading and interpreting the map's probabilistic indication could be difficult. For this reason, here we propose a simple yet informative alternative to conveying the full probabilistic prediction, in the form of mean values and uncertainty estimates around those. Our approach is to perform a first post-processing step where the mean posterior estimates are binned into a few classes. Here, we choose three for simplicity, to be plotted according to a standard traffic light criterion, corresponding to low (green), medium (yellow) and high (red) susceptibility values. We export all SUs belonging to a given class in a separate file, reporting the width of the 95% credible interval (CI) for each unit. We then plot each file separately, assigning a monochromatic colorbar whose intense colours correspond to SUs with low uncertainty and faded colours for SUs with high uncertainty. By plotting the three files together, we produce a unified probabilistic overview of the model in map form. We believe this to be a solution to a common problem between scientists and policymakers, and further consideration on this topic will be provided in Section 3.6.

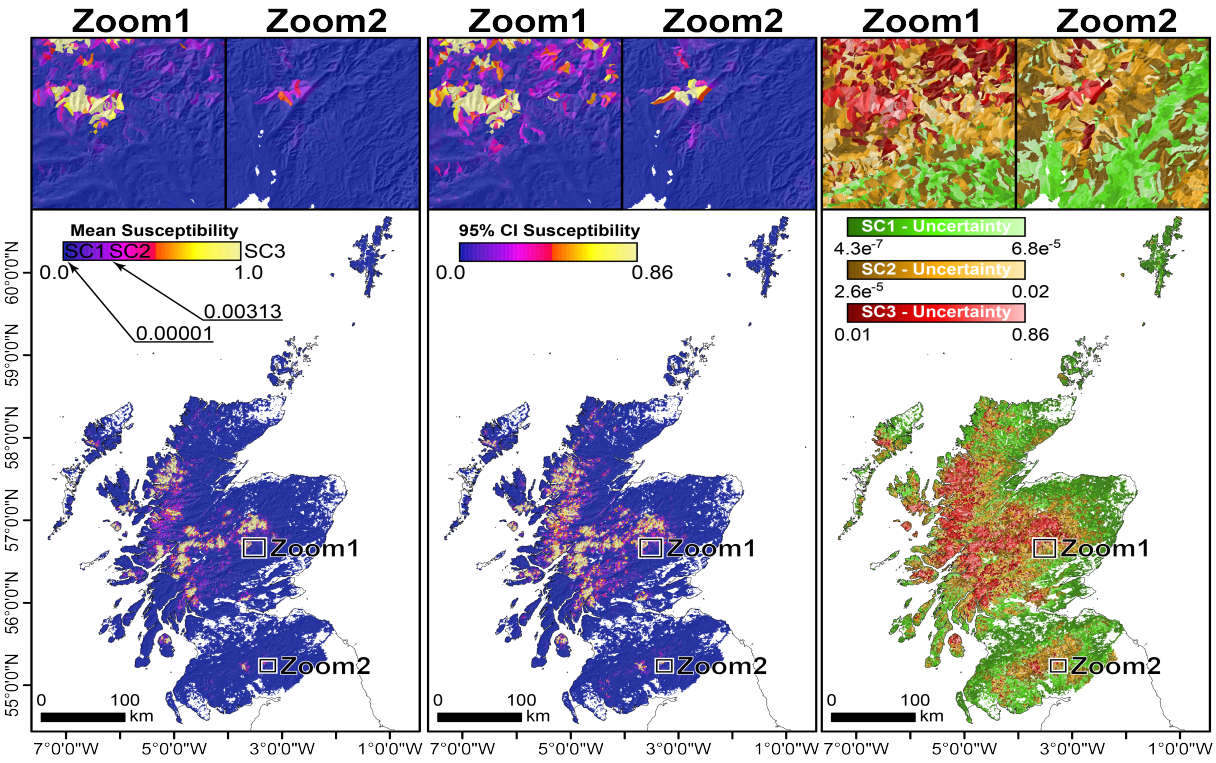

Figure 3.4: Mean DF susceptibility based on the presence/absence observations (left), the associated width of 95% CI (middle), and the combined mean susceptibility graded by its uncertainty (right).

3.4.2 LGCP

The marginal effects for the variables selected and shown in Table 3.3 are presented in Figure 3.5. Interestingly, the selected covariates boil down to those of the same nature as those selected in the susceptibility case. Specifically, meteorological and terrain characteristics control the variation of the DF intensity. However, the situation for the LGCP model is flipped as compared to the susceptibility one. In fact, the annual rainfall maxima over a period of 20 years appear to be much more closely linked to the response variable and morphometric characteristics that come after it. We recall once more that the intensity of an LGCP model can be considered a rate of DF occurrences in a given neighbourhood. Therefore, the spatial information this parameter conveys is ideally more complex than the simpler binary case tackled in a susceptibility task. This may be partially the reason behind the dominant contribution of rainfall extremes (the maximum among the yearly sums taken over a period of 20 years) towards the intensity. We would like to stress once more something partially referred to in the susceptibility case. Such a rainfall covariate should not be interpreted in the same way as in landslide early warnings (Guzzetti *et al.*, 2020), where the rainfall is measured in a much

narrower time window, comparable to the landslide failure process. In fact, the DF inventory we used does not report the landslide date and time of occurrence, therefore hindering the possibility of building temporal or spatiotemporal models. Conversely, the DF data representation we use is purely spatial and thus the rainfall interpretation needs to be simplified and summarised to a spatial context only. As a result, the maximum values among the annual accumulated ones over two decades exclusively reflect the geographic tendency of the Scottish landscape to be exposed to intense precipitation. Going back to the estimated effect, the first panel of Figure 5 highlights a gradual positive trend, distinguishable into two main near-linear segments. The first one starts at an approximate maximum of 1,000mm and continues with a similar incremental rate up to 2,100mm. Up to this point, the contribution still appears to decrease the expected rate of DF occurrences per SU. Conversely, from this point to around 2,500mm the effect shifts to a positive contribution to the estimated DF intensity, after which, it reaches a sort of plateau up to 4,500mm. The second largest contributor to the landslide intensity appears to be SLO μ. Different from the susceptibility case, here the mean SU steepness appears to be much more relevant, behaving according to a marked non-linear trend. This time, the effect is negative overall up to ten degrees, while showing a positive incremental trend that continues until 27 degrees. The third panel of Figure 3.5 shows the non-linear effect of local relief (LR) on DF intensity, which, as mentioned before, is a diagnostic of higher energy potential. The effect is negative up to 200m before increasing its positive effect until a spike at around 250m. From here, the effect is shortly negative in its influence before regaining a positive trend until 600m. The final panel of Figure 3.5 shows the linear effects of the standard deviation of the profile curvature and the standard deviation of the slope per SU, with profile curvature being negative ($\beta_{\text{PRC}\sigma} = -0.319$) and significant and the latter being positive ($\beta_{\text{SLO}\sigma} = 0.520$) and significant. The variation in the profile curvature per slope is something that we can interpret in terms of roughness. In other words, large variations would imply a rough terrain where the curvature measured across the vertical direction changes frequently in a stepped-like manner. Conversely, low variations would imply a relatively smooth surface. For this reason, we interpret a negative regression coefficient as an additional topographic information to that carried by SLO μ. Specifically, SUs with high average steepness values but with the same being largely kept constant across the vertical profile are prone to host large numbers of DFs. A justification for this may be assumed in the geotechnical interpretation of large curvature variations. In fact, mostly hard materials have the capacity to produce stepped-like landscapes, and they are mostly not capable of releasing shallow landslides but rather rockfalls or topples (Frattini and Crosta, 2013). As for constant- or near-constant steep slopes, these are diagnostic

of soft unconsolidated materials or soils draping over the bedrock. Thus, these are naturally the ideal hosts for DFs (Iverson, 1997).

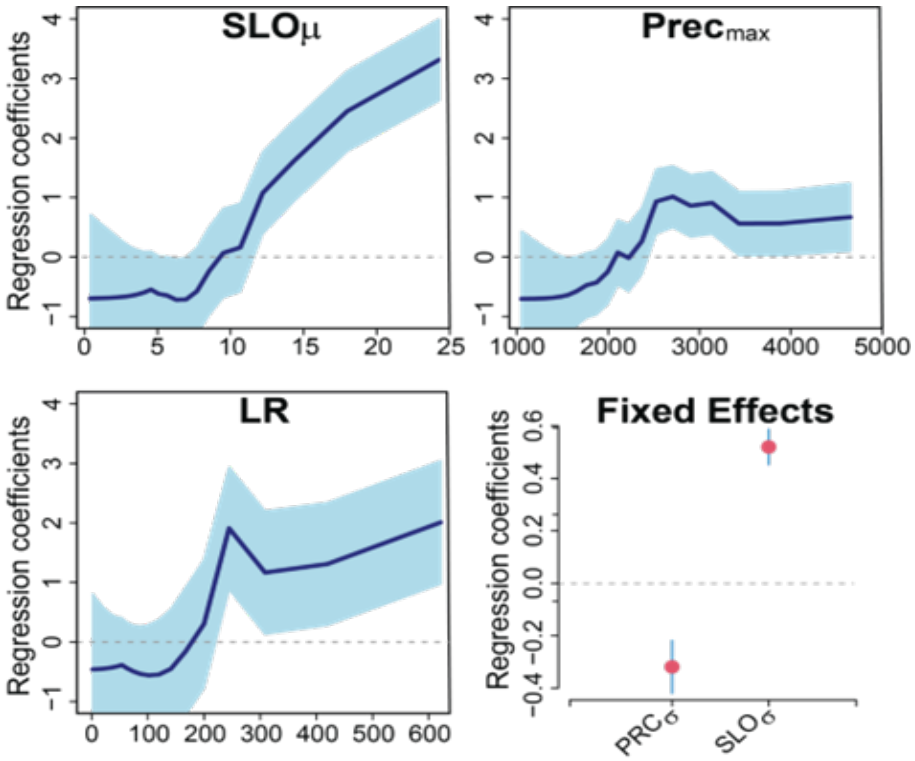

Figure 3.5: Posterior means (dark blue curves) of random effects with 95% pointwise CI (light blue shaded area) for the LGCP (panel a, b and c). Posterior means (dots) of fixed linear effects (except the intercept) with 95% CI (vertical segments) for the LGCP (panel d). The horizontal grey dashed lines indicate no contribution to the DF occurrence.

Analogously to the susceptibility case, the sum of all mean LGCP model components, together with the global intercept and, after the exponential transformation – required to convert the intensity from the log to the linear scale – produce the estimated mean DF intensity map shown in the left panel of Figure 3.6. What stands out is that the DF intensity is mostly concentrated on the West Coast and central Scotland. Something to be stressed here is that the patterns arising between the susceptibility and the intensity are quite similar. This being said, the information contained in the two maps is not the same. In fact, the susceptibility purely contains information on the occurrence probability whereas the intensity contains information closer to the requirement of hazard modelling. In fact, if we assume a mean DF size (area or volume), then a higher rate of DFs per SU would lead to a higher expected hazard. Conversely, if we consider an average DF size in the context of

susceptibility, the associated map will not account for the number of events, and therefore, for the expected hazard in a given SU. Another interesting element in maps shown in Figure 3.6 corresponds to the variability in the mean intensity estimates, shown in the central panel. What we see here is that the variation is minimal. This is comforting information because it generally indicates that the expected intensity or hazard associated with a given SU is robust. As for the last panel, the similarity that characterised susceptibility and intensity in their respective first two maps ceases to hold here. In fact, the pattern of the combined intensity/uncertainty highlights has less variability than what is shown for the susceptibility. This, in turn, may indicate that not only is the intensity mapping more informative than its susceptibility counterpart, but that the model is also more certain about its output. As a result, an ideal use of such a map may be more effective for decision makers.

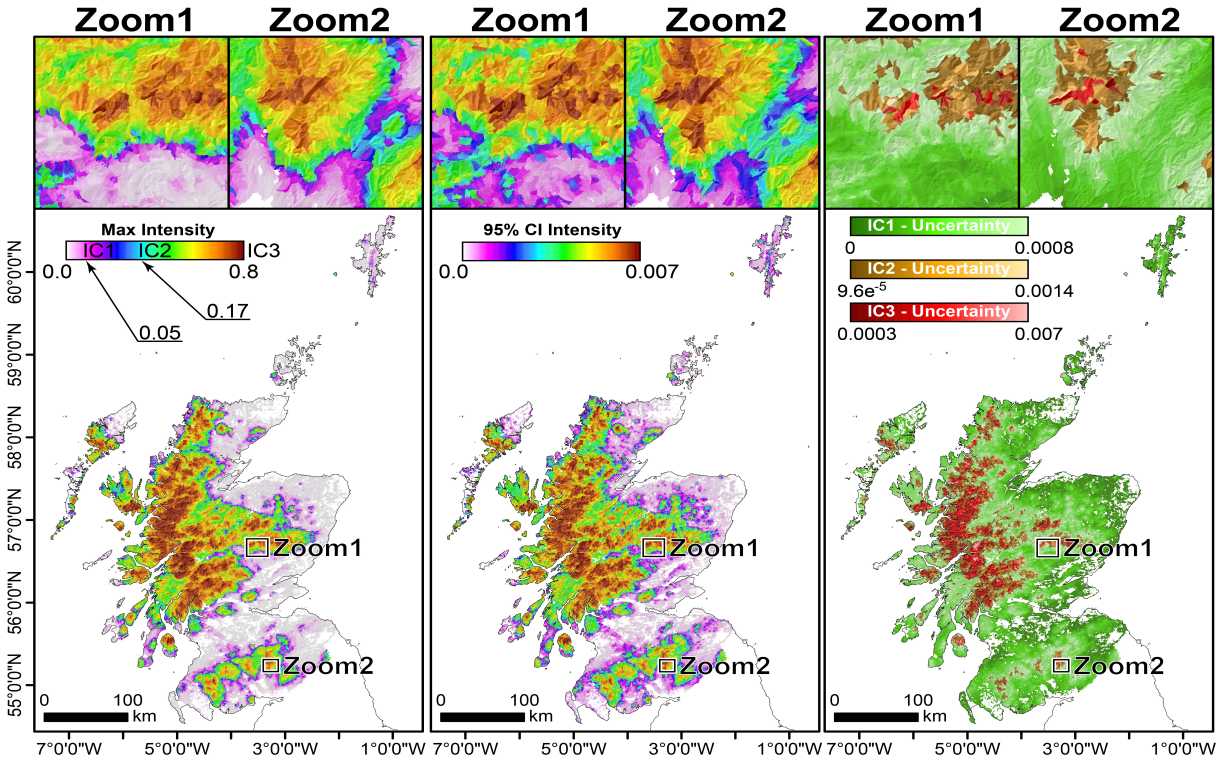

Figure 3.6: Mean DF intensity (left), the associated width of 95% credible intervals (middle), and the combined mean intensity graded by its uncertainty (right) aggregated across the Scottish SUs.

3.5 Model performance

In this section, an overview of the model performance is provided, spanning the fitting and cross-validation routines implemented for the susceptibility and LGCP models, respectively. The left panel

of Figure 3.7 shows the goodness-of-fit for the susceptibility case, through a ROC curve with an AUC of 0.97. We recall here that the ROC curve is a measure of the true positive rate (unstable SUs predicted by the model to be unstable, as a fraction of the total number of unstable SUs) against the false positive rate (stable SUs predicted by the model to be unstable, as a fraction of the total number of stable SUs) (Bewick *et al.*, 2004). Such a value is an indication of an extremely high explanatory power (outstanding according to the classification proposed by Hosmer (2000). To test whether this is due to overfitting, I implemented a 10-fold cross-validation (CV). This procedure involves sub-setting the dataset into ten random portions, each one made of 10% of all data. As a result, the union of the ten subsets returns the whole Scottish territory. The 10-fold prediction skill is graphically shown in the right panel, where the ten receiver operating characteristic (ROC) curves appear to showcase a limited, if not negligible, spread. This attests to the model's robustness. Aside from pure modelling considerations, going back to the susceptibility map, such high predictive power reflects the ability of the model to constrain the unstable labelling to the yellow region highlighted in Figure 4, where essentially most of the Scottish SUs that host at least one DF take place.

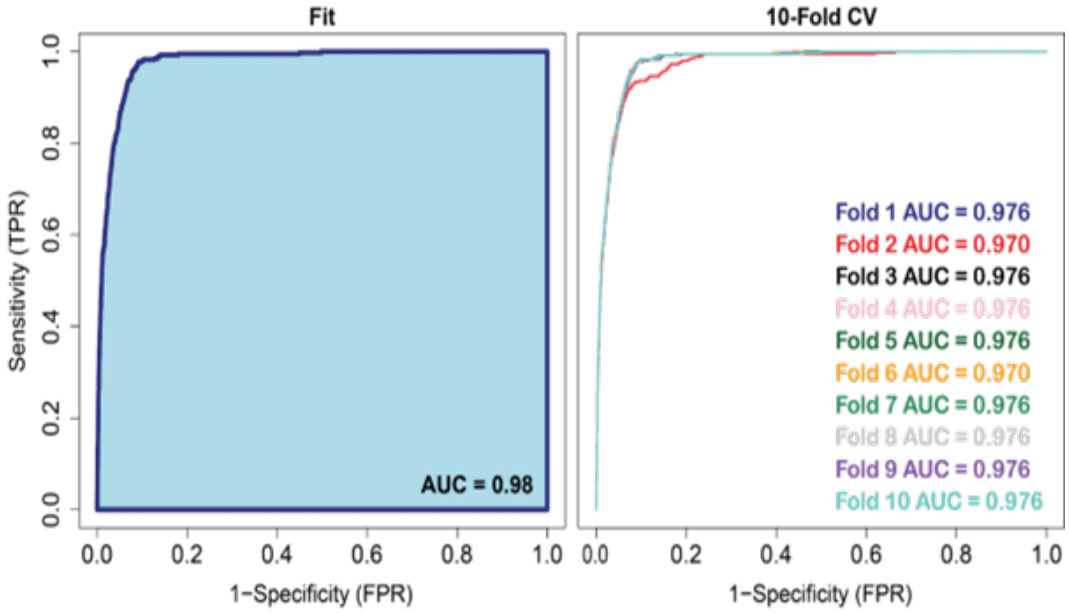

Figure 3.7: ROC curve and AUC value for the susceptibility model fit (left panel) and the 10-fold cross validation ROC curves and associated AUC values for the susceptibility model (right panel).

Validation and model assessment are generally complex for LGCP models, as the interest is in the position of the points in space rather than some value at said point. However, as the Poisson approximation to the LGCP likelihood was used, the same 10-fold cross-validation technique can

be used to divide the SUs and examine the DF counts as a function of the model's resulting mean intensity, thus creating an approximation for cross-validation for the LGCP. Looking at the results of the assessment counterpart for the LGCP, the performance also appears good, albeit less outstanding. The left panel of Figure 3.8 shows the observed count per SU versus the fitted count per SU, obtained by multiplying the SU mean intensity by the corresponding SU area. The agreement between the two parameters appears to hold for small counts. However, it shows an increasing deviance for large counts. This is most likely because very high counts are much less represented in the model. Therefore, small to medium counts are relatively easy to model because the LGCP learns from the available data. As for the medium to high counts, their small numbers hinder the ability of the LGCP to reflect them in the fitted results. The prediction skill of the LGCP is presented in the right panel of Figure 3.8, where the observed counts are plotted against their predicted counterparts in the same 10-fold division manner as the susceptibility case. The plot shows a similar behaviour as compared to the fit, with low to medium counts being suitably estimated. However, the prediction of medium to high counts is not as good. This attests once more to the model's robustness, where little variations are experienced when changing the modelled data. Similar to the susceptibility case, if we look at the locations where high mean DF intensity is shown in Figure 3.6, a high predictive power means that the region from SW to NE where high DF rates per SU are observed is consistently recognised.

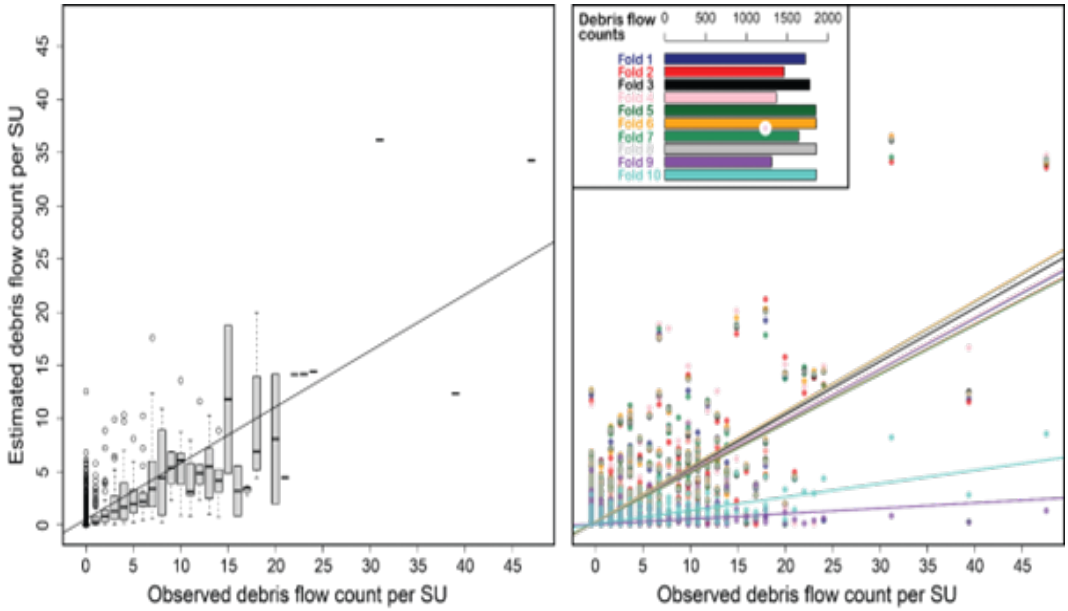

Figure 3.8: Observed versus fitted counts for LGCP model (left panel) and observed versus 10-fold predicted counts for LGCP model (right panel).

3.6 Discussion

Here we reflect on the noticeable effects of using one modelling procedure (susceptibility case) and the other (intensity case). The susceptibility model produced exceptionally high goodness-of-fit and predictive performance diagnostics (see Figure 3.7). When we first observed the outstanding classification (Hosmer, 2000) in the left panel, we assumed it to be potentially due to some clustering or spatial structure within the covariates. If that were the case, implementing a spatial cross-validation routine would have been capable of breaking down or reducing the contribution of any spatial structure, thus producing spatially unbiased predictive performances. As a result, we would have also noticed a marked decrease in the estimated classification metrics. However, the right panel still shows outstanding predictions. For this reason, we further investigated whether this could be due to some potentially biased covariates, as per Steger *et al.* (2021). To do so, we circled back and generated five single-variable models, one for each covariate selected in our benchmark susceptibility. Interestingly, extremely high performance is obtained solely using LR or SLO μ. These are not covariates that should be sensitive to any mapping criterion. In other words, when looking for biasing covariates, one should expect a property to explain a large portion of the DF distribution, this being the case because the covariate itself may be sensitive to the way local geological surveys are carried out to report landslides. For instance, Moreno *et al.* (2024) found that the effectively surveyed area (a layer expressing proximity to road networks) correlates well with the presence/absence of landslide data in South Tyrol (Italy), a bias they removed from the model by zeroing out the associated regression coefficient. In our case, we have no reason to assume that the inclusion of the LR and SLO μ is closely associated with any mapping practice behind the Scottish DF inventory. Despite this being an uncommon result, we believe our model to be reliable and the effect of these two terrain characteristics to be realistic. We support this argument by benchmarking our model against the GeoSure model, and by comparing the covariate effects estimated as part of the intensity procedure. Analogous susceptibility patterns can, in fact, be seen also in the GeoSure map (Figure 3.9). At the time of the GeoSure heuristic development, most of the DFs we used were not available, especially those that had been mapped in response to public notice. Therefore, it is highly unlikely that the public would report DFs, depending on the terrain arrangement, and it is rather more reasonable to assume that the slope geometry may largely contribute to the genesis of DF in Scotland. An additional verification can be seen in the intensity model. There, a more reasonable performance is obtained, far from being outstanding. This actually brings another point of discussion. After many

years of data-driven methodological development, outstanding performances have become commonplace among many susceptibility contributions. For instance, outstanding performance diagnostics, e.g. $AUC > 0.95$, are nowadays reported frequently in a number of articles adopting advanced spatial statistics (**?**), machine (Di Napoli *et al.*, 2020) and deep (Lv *et al.*, 2022) learning. Therefore, the point we are trying to raise here is questioning whether the susceptibility framework shouldn't be considered largely solved (Ozturk *et al.*, 2021), whenever heavily non-linear models are tasked with distinguishing the distribution of landslides purely in space. Conversely, the data-driven estimation of landslide intensity (Lari *et al.*, 2014), whether it is spatially (Moreno *et al.*, 2024), temporally (Nava *et al.*, 2023), or spatiotemporally (Fang *et al.*, 2024) addressed it is still at an infancy stage where few contributions are available and much may still be gained from a common geoscientific effort.

Here, the GeoSure DF susceptibility layer against the DF susceptibility produced are compared. Figure 3.9 allows for such comparison by reporting in panel (a) the original GeoSure raster at 50m resolution. Because our model is expressed at the SU scale, we opted to aggregate the 50m information over SUs, following two separate criteria: the first (panel (b)) assigns the most frequent DF susceptibility class, while the second (panel (c)) assigns the worst-case scenario. In other words, the first criterion calculates how many 50m pixels fall among the five GeoSure categories (VL = very low, L = Low, M = medium, H = high, VH = very high), and assigns to each slope unit the most representative class. The second approach assigns the highest class, irrespective of its numerosity per SU. We recall here that the definition of susceptibility refers to the probability of having at least one landslide occur. For this reason, we also included the second approach to account for at least one high probability pixel. What immediately stands out is that the probability patterns of our mean DF susceptibility and the pixel-based GeoSure largely match across Scotland. The same consideration applies when looking at the GeoSure majority aggregation per SU, but it is not the case for the GeoSure worst-case scenario. As for the match at the local level, GeoSure tends to polarise the susceptibility estimates either in the VL to L classes or in the H to VH classes, leaving the intermediate class less represented. This is mostly an effect due to the expert-based structure behind the current GeoSure maps, and it becomes particularly evident in the Zoom1 of panel (a), where the effect of the slope steepness map largely controls the susceptibility classes, with no other apparent contribution coming from other predictors. The same is true for Zoom2 of the same panel. There, the effect of a geological type that has received a negative weight flattens out the susceptibility, which mostly falls in the VL category. By comparison, our mean DF susceptibility provides a richer description of the process, not only because it includes the uncertainty, but also because its patterns appear more realistic. Another

interesting element to consider has to do with the pixel structure of a susceptibility map, and this reasoning goes well beyond the specific case and rather applies to any similar gridded map versus a SU-based one. In fact, if one selects grids as the mapping unit of choice, the most common effect is to obtain "salt and pepper" looking maps. This is again visible in Figure 9 panel (a), where the zooms highlight single pixels falling in the VH susceptibility class being entirely surrounded by pixels falling in the L or even VL classes. This constitutes a problem for decision-makers because slope stabilisation practices cannot be applied to single grid cells but rather to the whole slope they belong to. This 'salt and pepper' effect propagates to the susceptibility in panel (c). Risk-oriented applications are often tailored towards worst-case scenarios, and here, what becomes evident is that almost every SU in Scotland hosts at least one VH susceptibility pixel. However, not every slope in Scotland fails and luckily so, which in turn points to the limited realistic representation of a landscape in grid-based models.

Firstly, we note that the data is as complete and representative as possible. Therefore, modelling the higher DF counts is unlikely to be improved with this methodology as there is no way to gain more data on higher counts if none exist. However, we can extend the data framework to account for a larger domain - the whole of GB, for example - and in this way we can gather more information on the spread of DF count and its dependence on the covariate set we chose.

Secondly, the covariate information and the modelling methodology are purely spatial. This is an informative start, but extending the point process modelling towards spatio-temporal structures could explain varying patterns of DF intensity. However, some degree of variation to the model should be implemented. For instance, the covariates should be revisited. Presently, we use an aggregate of the maximum daily precipitation over a 20-year period, calculating averages on an annual basis. Rainfall has a high correlation with DF intensity, but with a spatial model, we can only account for one statistic (the average of the maxima) to describe the whole rainfall pattern. If we used a spatio-temporal LGCP model, we could use a function to describe the pattern of rainfall over a period of time that might influence the slope instability, prior to and past the DF event. This would improve prediction ability and provide a model that is interpretable over time. In turn, this could open up towards a new generation of early warning systems for Scottish DFs. However, it should be stressed that not all the data points we used have an associated date, and ideally, we would want the complete data. The geoscientific community is working hard to improve this, mainly in the form of automated mapping procedures, thanks to the high orbital frequency modern satellites offer.

Overall, the DF susceptibility and DF intensity maps both capture the areas in which to focus

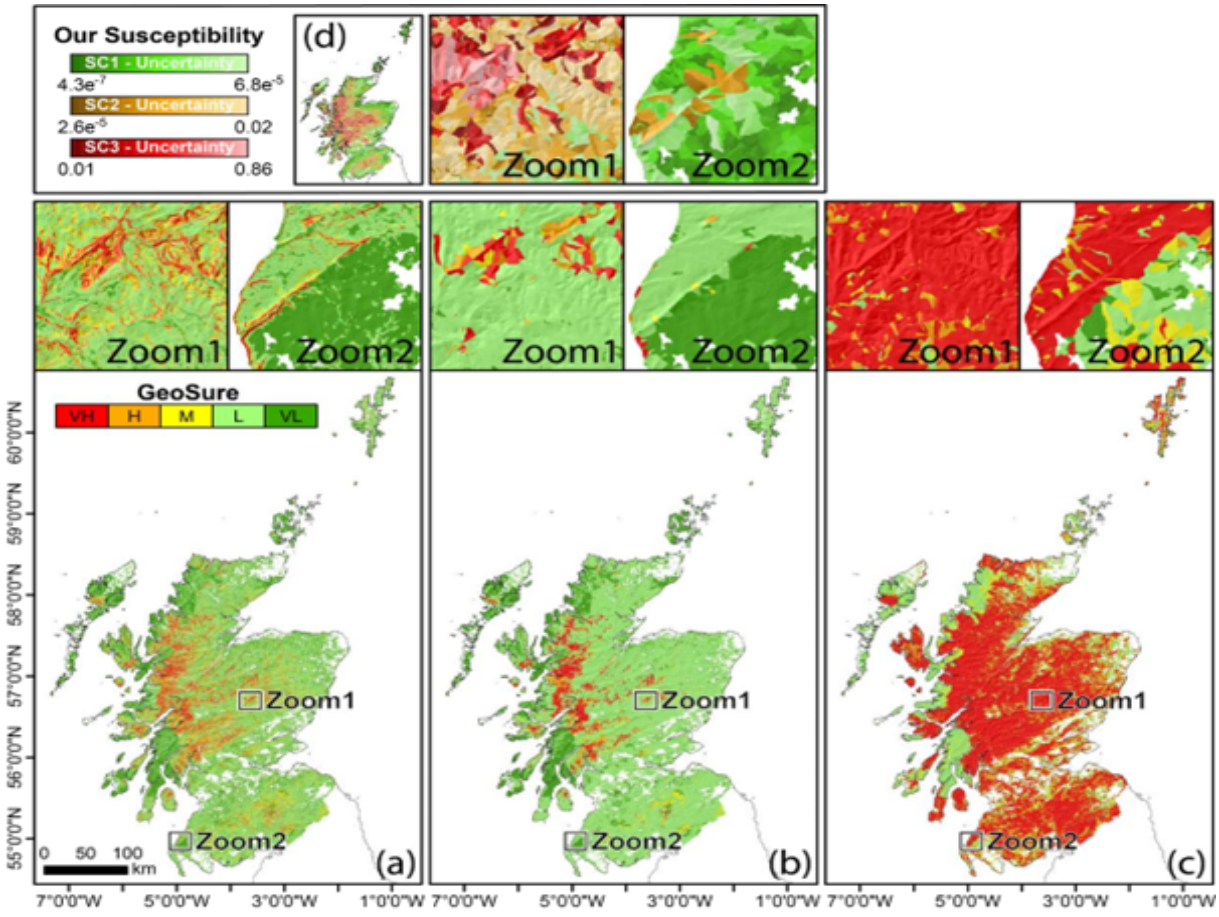

Figure 3.9: GeoSure DF susceptibility (panel a), aggregated on the basis of a majority criterion per SU (panel b), and according to the highest class per SU (panel c). Panel d shows a reduced version of the DF susceptibility built in this work and already mapped in Figure 3.4.

in terms of a higher DF risk. The LGCP model intensity map, however, perhaps pinpoints these areas with a higher degree of accuracy due to the nature of the point process modelling approach. Both models do well in terms of model performance, although validation measures for point-process models are generally complex and more along the lines of a residual analysis to compare variations of the model. Using the DF count per SU allowed us to implement the same 10-fold CV scheme that was used to validate the susceptibility model. An improvement here could be conducting a leave-one-group-out cross-validation (Liu and Rue, 2022). When one SU is removed, the underlying spatial correlation between SUs can still be closely approximated by the surrounding SUs. Removing a group of SUs at a time would better test the model's prediction abilities by accounting for the absence of this spatial correlation when a group of SUs are removed.

A final improvement to the model can be achieved with the integration of information from the

BGS National Geotechnical Database on the geographical distribution of physical properties of a wide range of rocks and soils present in GB. Presently, the information is relatively coarse across space, which makes its integration into the model difficult at this time. BGS is continuously updating its records and databases, which in the future could be used in a model for DF prediction. We conclude the discussion by pointing out an interesting aspect in the landslide predictive patterns coming from our susceptibility and intensity models, as well as the susceptibility produced for the GeoSure system. The general patterns of the three corresponding maps look similar at a broad examination. The main areas where landslides are expected follow two trends, one North to South on the West coast, and another one South-West to North-East starting from the West coast and extending across the central belt. This is interesting because both our models rely on underlying spatial effects, whereas the GeoSure map is purely driven by an expert-based weighting system informed by knowledge of the debris flow initiation process.

3.7 Challenges in modelling debris flow intensity with LGCPs

A key issue emerged during the LGCP extension: in the initial implementations, the spatial intensity of the point pattern was almost entirely absorbed by the SPDE-based random effect. As a result, the inclusion of environmental covariates had little to no influence on the fitted intensity, which exhibited an unrealistic block-like pattern (see right panel of Figure 3.10). To address this in the context of the BGS project, we adopted a count-based approximation to the LGCP likelihood, as described in Section 3.3. However, this problem resurfaced in the subsequent spatio-temporal case study, where the same modelling structure was extended. Chapter 4 is therefore devoted to a detailed examination of this phenomenon—its causes, consequences, and possible strategies to mitigate its impact.

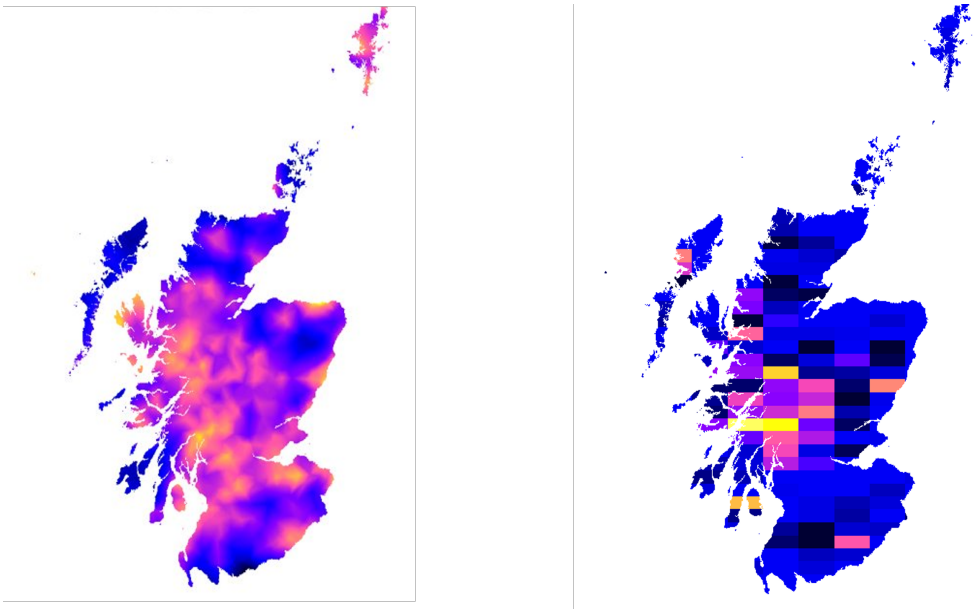

Figure 3.10: Debris flow intensity of LGCP model explained by random spatial effect only (left panel) and DF intensity of LGCP model explained by an environmental covariate only (right panel).

Chapter 4

Mess with the mesh: High-resolution covariates for point process models in INLA

This chapter investigates a methodological challenge that arises when implementing the SPDE approach in point process models using high-resolution covariates. The issue is first examined through a series of simulation studies, and a practical solution is presented through a case study where a marked point process model is fitted to a Japanese landslide inventory.

4.1 Introduction

In this chapter, we revisit a challenge encountered in the log-Gaussian Cox process (LGCP) extension for modelling debris flows across Scotland, outlined in Section 3.7. The problem is encountered specifically in point process models when the spatial resolution of the covariate information does not match, or varies rapidly between, the triangulation mesh defined for use with the stochastic partial differential equation (SPDE) approach.

Spatial point patterns are commonly characterised by an intensity function, defined as the expected number of events per unit area. This intensity is typically modelled as a function of spatial covariates to account for spatial heterogeneity, with any residual spatial structure captured via a spatial random effect. Within the SPDE approach, the latent Gaussian random field (GRF) is approximated by the solution to an SPDE, represented as a weighted linear combination of basis functions. These basis

functions are constructed to approximate the Matérn covariance structure (Matérn, 1986b), with the associated Gaussian weights indexed by the vertices of a triangulation mesh (see Section 1.5 for details). The mesh plays a critical role in discretising the continuous spatial domain, thereby enabling practical computation in models defined over continuous space.

Spatial covariates are similarly discretised by evaluating their values at specific integration points —-typically the mesh nodes—- and interpolating between these points to approximate their continuous spatial pattern. This ensures consistency between the representation of the spatial field and the covariates across the domain. However, this discretisation introduces a critical issue: when covariates exhibit rapid spatial variation, the interpolation may poorly represent their true behaviour, and different choices of integration points can lead to substantially different covariate values being included in the model.

The above highlights a broader concern—frequently misunderstood—about the relationship between mesh resolution and parameter estimates in spatial models (see, e.g., Dambly *et al.* 2023, Dovers *et al.* 2024). While the mesh specification must be carefully defined, it should not exert enough influence to fundamentally alter the estimated effect of covariates. In this chapter, we formally explore the problem observed in our LGCP debris flows model, generalise it through simulation, and present a case study application where we successfully implement a marked point process model using high-resolution covariate data. Our aim is to clarify misconceptions surrounding the role of the mesh in SPDE-based point process models and to raise awareness of the restrictions of combining fine-scale explanatory data with spatial discretisations.

4.2 Problem outline

This section clarifies the implications of evaluating both the covariates and the Gaussian Random Field (GRF) at a finite set of integration points, as is standard in spatial point process modelling. A key underlying assumption is that these integration points adequately represent the spatial variation of covariates and latent effects across the entire study domain. Thus, the validity of model inference depends crucially on the spatial resolution and distribution of these points.

We consider a spatial point process model of the form:

$$\begin{aligned} N(A) &\sim \text{Poisson}\left(\int_A \lambda(\mathbf{s}) d\mathbf{s}\right) \\ \log(\lambda(\mathbf{s})) &= \eta(\mathbf{s}), \end{aligned} \tag{4.1}$$

where $\mathbf{s} = \{s_1, ..., s_n\}$ are the spatial units that discretise a study region, A, $\lambda(\mathbf{s})$ is the intensity of the point pattern, and $\eta(\mathbf{s})$ is a structured additive predictor:

$$\eta(\mathbf{s}) = \alpha + \sum_{m=1}^{M} \beta_m w_m(\mathbf{s}) + \sum_{k=1}^{K} f_k(z_k(\mathbf{s})) + \boldsymbol{\omega}(\mathbf{s}), \tag{4.2}$$

where α is an intercept and $(w_1(\mathbf{s}), ..., w_M(\mathbf{s}))^T$ are covariates that might influence, in a linear way, the variable of interest, with coefficients $\boldsymbol{\beta} = (\beta_1, ..., \beta_M)^T$. The functions $f = \{f_k(\cdot), ..., f_K(\cdot)\}$ are non-linear effects defined in terms of covariates $(z_1(\mathbf{s}), ..., z_K(\mathbf{s}))^T$. The specific form of the functions $f_k(\cdot)$ can be those of a Gaussian random walk of order 1 or 2, for example (RW1/RW2; Krainski *et al.* (2018)). Finally, the term $\boldsymbol{\omega}(\mathbf{s})$ is a spatially structured random effect governed by a GRF. The GRF is approximated using the SPDE formulation, with the solution expressed as:

$$\boldsymbol{\omega}(\mathbf{s}) = \sum_{i=1}^{n} \omega_i \phi_i(\mathbf{s}),$$

where $\{\phi_i(\mathbf{s})\}_{i=1}^{n}$ are a set of basis functions needed to approximate the Matérn covariance, ω_i are associated Gaussian weights and n is the number of vertices in the triangulation mesh. Denoting $Y = N(A)$, the likelihood of the model in (4.1) is

$$\pi(\mathrm{Y}|\lambda) = \exp\{|A| - \int_A \lambda(\mathbf{s}) d\mathbf{s}\} \prod_{s_i \in \mathrm{Y}} \lambda(s_i).$$

Therefore, the exponent term is evaluated on the domain while the product term is evaluated at the observed points of the process. The corresponding log-likelihood is defined as

$$\begin{aligned}
\log \pi(\mathrm{Y}|\lambda) &= |A| - \int_A \lambda(\mathbf{s}) d\mathbf{s} + \sum_{s_i \in \mathrm{Y}} \log \lambda(s_i) \\
&= C + \int_A \exp\left\{\beta_0 + \beta^T \mathbf{X}(\mathbf{s}) + \boldsymbol{\omega}(\mathbf{s})\right\} d\mathbf{s} + \sum_{s_i \in \mathrm{Y}} \left\{\beta_0 + \beta^T \mathbf{X}(s_i) + \omega(s_i)\right\} \\
&\approx C + \int_A \exp\left\{\beta_0 + \beta^T \mathbf{X}(\mathbf{s}) + \sum_{j=1}^{n} \omega_j \phi_j(\mathbf{s})\right\} d\mathbf{s} + \sum_{s_i \in \mathrm{Y}} \left\{\beta_0 + \beta^T \mathbf{X}(s_i) + \sum_{j=1}^{n} \omega_j \phi_j(s_i)\right\},
\end{aligned}$$

where the integral over the domain has been approximated using a quadrature rule:

$$\int_A f(\mathbf{s}) d\mathbf{s} \approx \sum_{i=1}^{p} \tilde{\alpha}_i f(\tilde{s}_i)$$

with $\tilde{s}_i$ denoting the integration points and $\tilde{\alpha}_i$ their associated weights. Therefore, we have

$$\int_A \exp\left\{\beta_0 + \beta^T \mathbf{X}(\mathbf{s}) + \sum_{j=1}^{n} \omega_j \phi_j(\mathbf{s})\right\} d\mathbf{s} \approx \sum_{i=1}^{p} \tilde{\alpha}_i \exp\left\{\beta_0 + \beta^T \mathbf{X}(\tilde{s}_i) + \sum_{j=1}^{n} \omega_j \phi_j(\tilde{s}_j)\right\}.$$

This approximation highlights the fundamental modelling assumption often not mentioned explicitly, which is that the spatial variability of both the covariates and the latent field is sufficiently captured by the values at a finite number of integration points. In practice, this raises important considerations for model design:

- If integration points are too sparse or poorly placed, key covariate patterns or spatial dependencies may be missed, leading to biased or imprecise inference.
- Increasing the number of integration points can improve approximation accuracy, but at the cost of higher computational complexity and slower model fitting.

The mismatch between covariate resolution and the mesh leads to biased estimates because the integration scheme may not capture fine-scale covariate variation accurately. This can distort the estimated effect sizes and compromise inference (imprecision), especially for covariates with short spatial ranges. Thus, ensuring alignment between mesh resolution and covariate structure is crucial for valid interpretation. This trade-off is particularly critical in applications where covariates exhibit fine-scale variability, or where spatial units (e.g., slope units or hydrological catchments) are defined based on criteria relevant to the system being modelled, e.g. geomorphological criteria as in the example discussed here. While it is tempting to increase spatial resolution to capture local detail and resolve the issue, doing so may compromise computational feasibility, especially in Bayesian frameworks with latent Gaussian models. A careful balance between spatial accuracy and computational tractability must therefore be maintained.

In the following sections, we explore the impact of integration point resolution on model estimates and propose guidance for effective mesh and integration scheme design.

4.3 Simulation study 1: sensitivity to mesh alignment and covariate range

This simulation study investigates how the positioning of the triangulation mesh affects parameter estimates in spatial point process models, particularly in relation to the spatial correlation range of covariates. The aim is to assess the robustness of model estimates when the mesh domain is shifted relative to a fixed spatial grid.

We define six triangulation meshes over a regular grid covering the domain $(-100, -50, \ldots, 550, 600) \times (-100, -50, \ldots, 550, 600)$. Each mesh is systematically shifted relative to the grid cells—either to the

left, right, above, or below—to examine the sensitivity of model outputs to the mesh domain. This set-up is illustrated in Figure 4.1. The left panel displays the six meshes, with red boundary lines highlighting the domain shifts applied to each configuration. The right panel visualises the corresponding integration domains, where differences across meshes are reflected in the spatial arrangement of the integration nodes, shown in distinct colours.

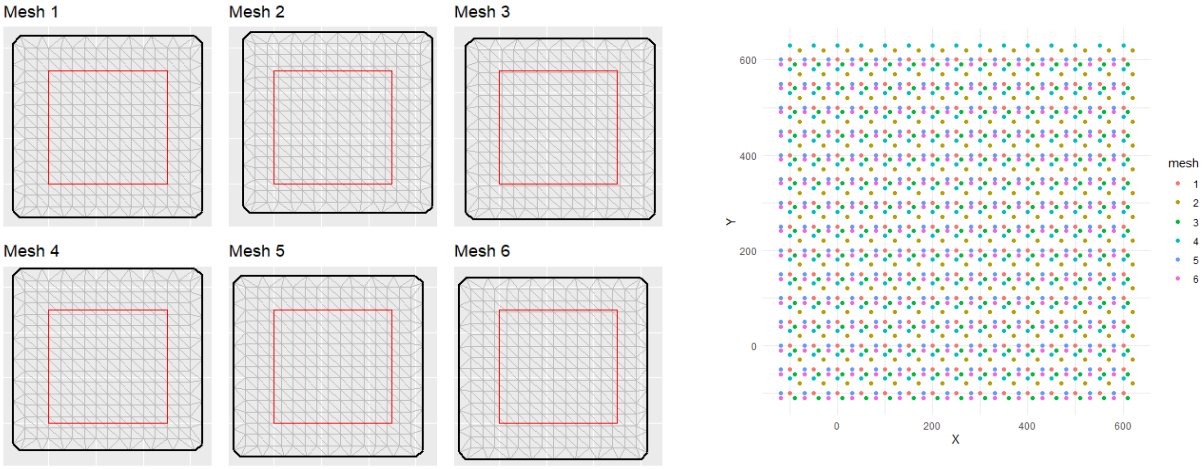

Figure 4.1: Simulation study 1: shifted meshes and location of each mesh node.

To explore how covariate smoothness interacts with mesh alignment, we simulate three covariates as Gaussian random fields with increasing spatial range: 10 (short-range), 50 (medium-range), and 600 (long-range). The covariates are simulated in the following way:

1. Define a 2D triangular mesh over the bounding polygon (square with corners from (0, 0) to (500, 500)) and `max.edge` = c(10,40) to discretise the spatial domain for the GRF. Note that this is a different mesh from the ones defined in Figure 4.1.

2. Define a Matérn SPDE model.

3. Define a projection matrix to map all integration points (mesh nodes and coordinates of covariate).

4. Define the precision matrix of the Matérn field.

5. Using the precision matrix, sample from a multivariate normal distribution through `inla.qsample`.

6. Project the sampled field across the integration points.

Once this is done, each covariate is then used to simulate a spatial point pattern with log-linear intensity function:

$$\log(\lambda_i) = \beta_0 + \beta_1 X_i,$$

where $\beta_0 = -6.5$, $\beta_1 = -0.8$, X_i is the $i-$th covariate, and $i = 1, 2, 3$; Figure 4.2 shows the resulting covariate fields and corresponding point patterns.

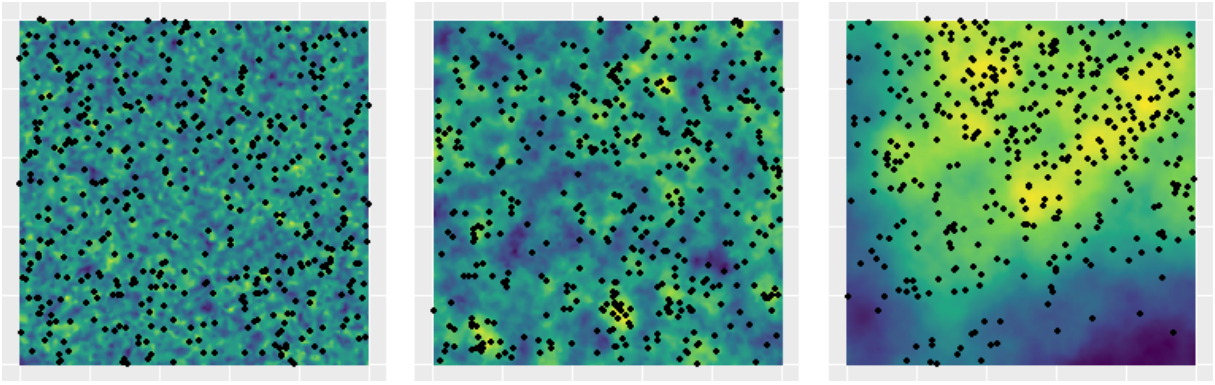

Figure 4.2: Simulation study 1: covariate fields and simulated point patterns for short (left), medium (middle), and long (right) spatial range.

We fit an inhomogeneous Poisson point process model to each of the three datasets using all six mesh configurations. Figure 4.3 summarises the results. The first column displays the estimated covariate effects, $\hat{\beta}_1$, while the second column displays the estimated intercepts, $\hat{\beta}_0$, with horizontal dashed lines indicating the true values. Although only the domain of integration is adjusted, notable variation is observed in the estimated covariate effects, particularly for Covariate 1, which has the smallest spatial range. Estimation accuracy improves moderately for Covariate 2 (intermediate range) and becomes most stable for Covariate 3 (largest spatial range). These results highlight the sensitivity of parameter estimation to the specification of the integration domain. Even minor shifts in the locations of the integration points—such as displacement by a single grid cell—can lead to clear differences in the estimated values and hence conclusions drawn from the modelling results. It is clear that in the estimation of the intercept, the shifting of the domain of integration plays a role (variability in estimation throughout first and second datasets). For the third dataset, the estimation of the intercept is relatively consistent for each mesh, although still less than the true value.

Interestingly, the most accurate estimates for both the covariate effect and intercept do not consistently correspond to the covariate with the largest spatial range (Covariate 3), contrary to what might be intuitively expected. This discrepancy is likely influenced by the interplay between the

spatial range of the covariate and the mesh structure itself—a relationship explored further in the subsequent simulation study.

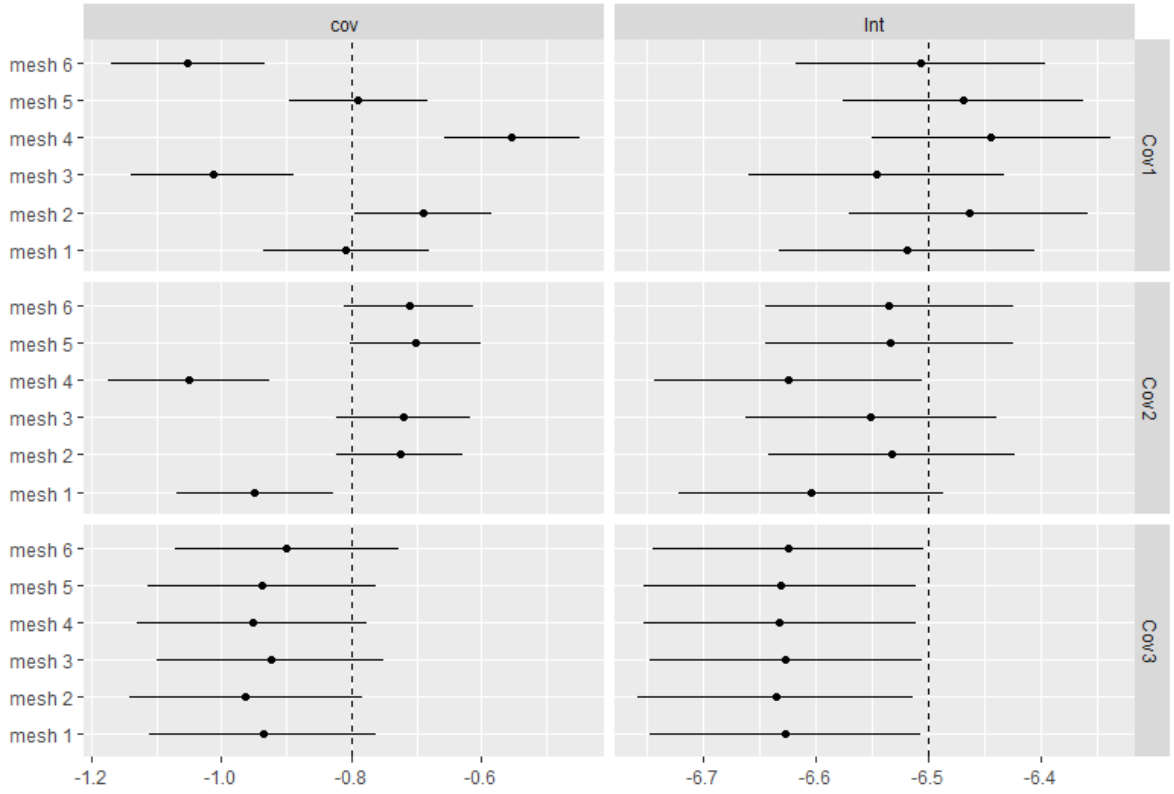

Figure 4.3: Simulation study 1: estimated covariate effects (left column) and intercepts (right column) across six mesh configurations, for each of the three spatial ranges. Dashed lines indicate true parameter values.

4.4 Simulation study 2: impact of mesh resolution on parameter recovery

This simulation study explores how triangulation mesh refinement influences parameter estimation in spatial point process models. In particular, we investigate how mesh resolution interacts with the spatial smoothness of covariates and the inclusion of a latent Gaussian random field (GRF) defined via a Matérn SPDE.

We define three meshes of increasing resolution using `fm_subdivide(mesh, n)`, with $n = 2$, yielding approximately $n_1 = 261, n_2 = 2233$ and $n_3 = 19765$ nodes. Figure 4.4 shows the resulting

triangulations. The refinement increases the number of integration points by subdividing each triangle into $(n+1)^2$ sub-triangles. We simulate three spatial covariates with increasing correlation range—10,

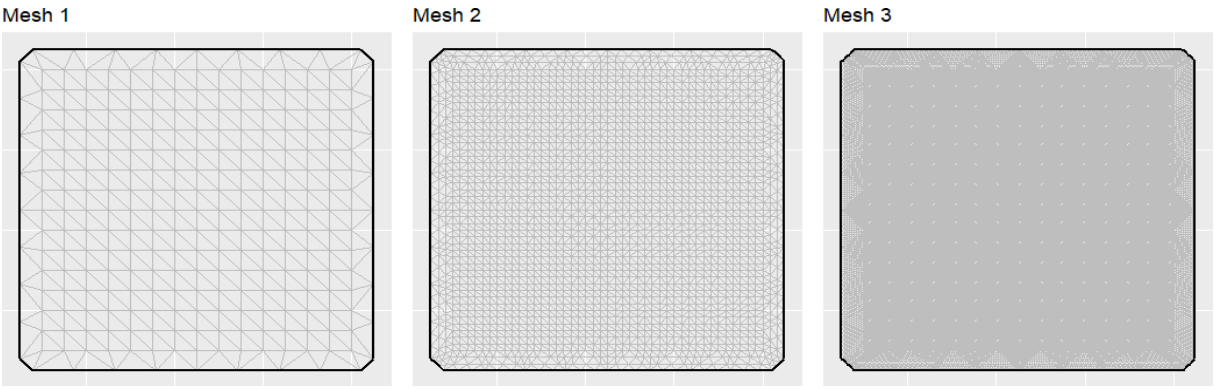

Figure 4.4: Simulation study 2: three triangulation meshes, increasing in resolution from left to right.

50, and 600 units—using `inla.qsample()`. The spatial random effect is represented by a realisation of a GRF, simulated in the same manner as the covariates, with a range of $r = 300$ and a standard deviation of $\sigma = 1$. The intensity of each point pattern corresponds to:

$$\log(\lambda_i) = \beta_0 + \beta_1 X_i + \omega_i,$$

where $\beta_0 = -7.5$, $\beta_1 = -1.3$, X_i is the $i-$th covariate, ω is the spatial random effect and $i = 1, 2, 3$. Figure 4.5 illustrates the three covariates (from small range to large range) with their associated point patterns and the GRF. Model fits using each of the three meshes are summarised in Figure 4.6. We

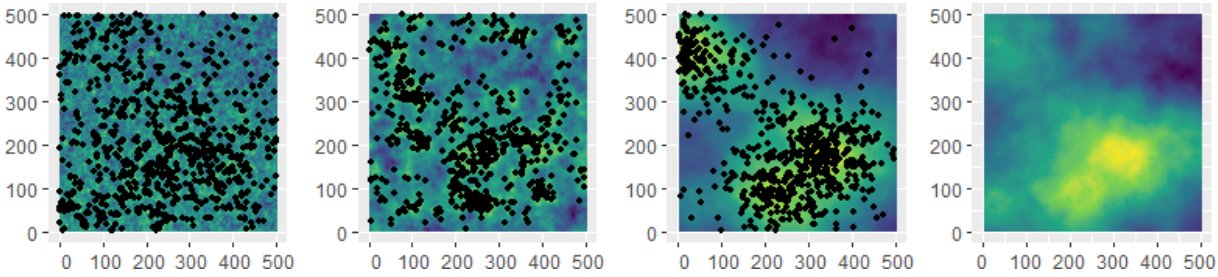

Figure 4.5: Simulation study 2: simulated point patterns, covariate fields and Matérn field.

examine estimates of β_1, β_0, and the GRF hyperparameters r and σ.

Covariate 1 (short spatial range): The coarser meshes (Mesh 1 and 2) yield wildly inaccurate posterior mean estimates of the regression coefficients and hyperparameters:

- Mesh 1: $\hat{\beta}_1 = -813.93$, $\hat{\beta}_0 = -329.87$, $\hat{r} = 2.07$, $\hat{\sigma} = 132.03$.

- Mesh 2: $\hat{\beta}_1 = -290.46$, $\hat{\beta}_0 = -96.45$, $\hat{r} = 1.12$, $\hat{\sigma} = 803.43$.
- Mesh 3: $\hat{\beta}_1 = -1.27$, $\hat{\beta}_0 = -6.84$, $\hat{r} = 319.16$, $\hat{\sigma} = 0.79$.

Covariate 2 (medium range):

- Mesh 1 still performs poorly: $\hat{\beta}_1 = -1392.12$,, $\hat{\beta}_0 = -403.58$, $\hat{r} = 34.00$, $\hat{\sigma} = 18.70$.
- Mesh 2 improves significantly: $\hat{\beta}_1 = -1.38$, $\hat{\beta}_0 = -6.93$, $\hat{r} = 453.00$, $\hat{\sigma} = 0.85$.
- Mesh 3 remains accurate: $\hat{\beta}_1 = -1.31$, $\hat{\beta}_0 = -6.89$, $\hat{r} = 444.05$, $\hat{\sigma} = 0.82$.

Interestingly, Mesh 3 slightly overestimates the range here, indicating that a finer mesh may not always yield the best result when the underlying field is smoother.

Covariate 3 (long range):

- Mesh 1 again breaks down: $\hat{\beta}_1 = -2438.42$,, $\hat{\beta}_0 = -46.97$, $\hat{r} = 0.87$, $\hat{\sigma} = 5279.51$.
- Meshes 2 and 3 show reduced accuracy compared to Covariate 2, possibly due to over-refinement:
  - Mesh 2: $\hat{\beta}_1 = -1.63$, $\hat{\beta}_0 = -6.83$, $\hat{r} = 658.62$, $\hat{\sigma} = 0.96$.
  - Mesh 3: $\hat{\beta}_1 = -1.56$, $\hat{\beta}_0 = -6.88$, $\hat{r} = 578.01$, $\hat{\sigma} = 0.95$.

This study shows that accurate estimation requires matching mesh resolution to the spatial scale of variation in both covariates and the underlying GRF. For fine-scale features, coarse meshes lead to biased estimates and unstable hyperparameter recovery. Over-refinement, on the other hand, can reduce stability for smoother fields. These findings emphasise the need for careful mesh design tailored to the spatial properties of the data. Specifically, this refers to the `prior.range`(ρ_0, p_ρ) argument when defining a SPDE which defines the probability $P(\rho < \rho_0) = p_\rho$, where ρ is the real spatial range of the random field, and the `prior.sigma`(σ_0, p_σ) argument which defines the probability $P(\sigma > \sigma_0) = p_\sigma$, where σ is the true marginal standard deviation of the field. These spatial properties alongside suitable mesh triangulation metrics (`max.edge, min.angle, cutoff, offset`) are imperative to match mesh to data.

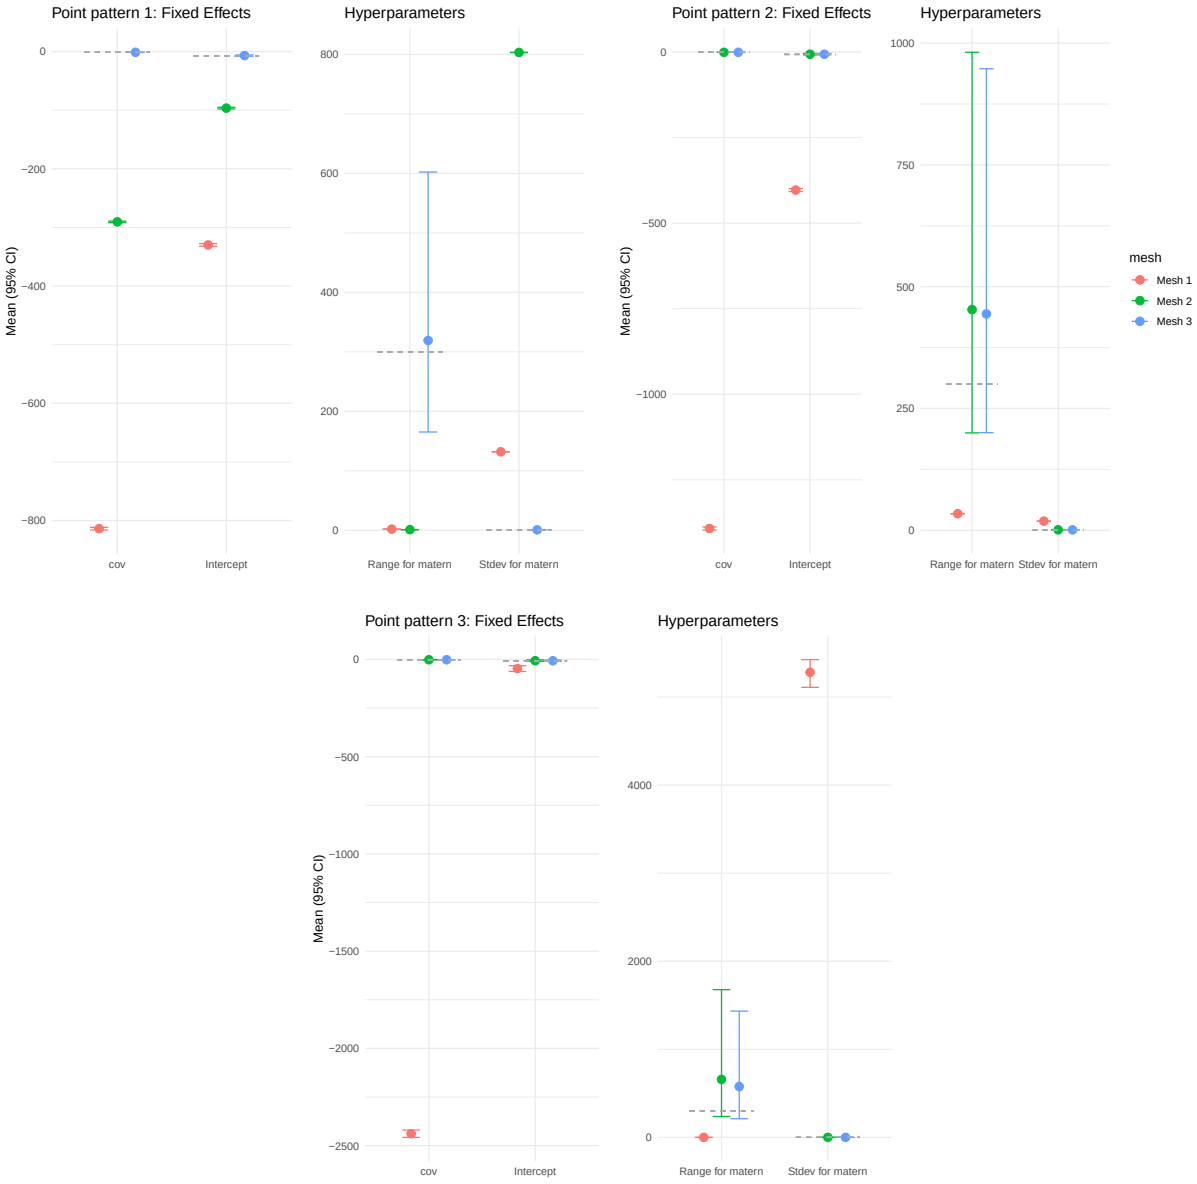

Figure 4.6: Simulation study 2: estimated parameters across three mesh resolutions for each simulated point pattern.

4.5 A case study: Japanese landslide inventory

This section presents a real-world case study demonstrating the application of a marked log-Gaussian Cox process (LGCP) to model spatial point pattern data using high-resolution environmental covariates in conjunction with the SPDE approach. Building on the considerations and simulation studies from earlier in the chapter, we apply these methods to a Japanese landslide inventory dataset, with

the goal of capturing both the spatial distribution and associated magnitudes of landslide events.

A key motivation for this analysis is to illustrate how high-resolution covariate data can be effectively integrated into a point process framework without allowing the mesh specification to unduly influence model inference. As discussed in the introduction to this chapter, careful attention must be paid to the role of the properties of the triangulation mesh in SPDE-based models, particularly when explanatory variables exhibit fine-scale spatial variation.

A marked point process extends a spatial point process by associating each point, i.e. each event location $\mathbf{s} \in \mathcal{S} \subset \mathbb{R}^2$ with an auxiliary variable or mark $m(\mathbf{s})$. Marks can be categorical or continuous and provide additional information about each event (Illian *et al.*, 2008). We have defined the LGCP in (4.1) and (4.2) provides the expression for the linear predictor. In a marked LGCP, the marks $m(\mathbf{s})$ are modelled conditionally on location. In our case, the marks represent landslide sizes, which we assume to follow a log-Skew-Normal distribution. That is,

$$m(\mathbf{s}) \mid \mathbf{s} \sim \text{log-Skew-Normal}(\mu(\mathbf{s}), \sigma^2, \nu) \quad \Leftrightarrow \quad \log(m(\mathbf{s})) \mid \mathbf{s} \sim \text{Skew-Normal}(\mu(\mathbf{s}), \sigma, \nu), \tag{4.3}$$

where σ^2 denotes the variance and ν controls the skewness. The location parameter $\mu(\mathbf{s})$ may depend on the same covariates and latent field used in the intensity component of the model, allowing for a coherent joint modelling of occurrence and size. This distribution allows asymmetric heavy-tailed behaviour in the log-marks, which is appropriate for landslide sizes. In this way, we can focus on two core components of landslide hazard assessment as defined by Guzzetti *et al.* (1999): the spatial distribution of landslide occurrences and their size.

Given the computational demands associated with fitting LGCPs to large, high-resolution datasets, we focus on regional subsets of the Japanese territory rather than the full domain. For each subset, we implement a manual integration scheme by refining the mesh to better align with the spatial resolution of the covariate data, which is provided at the catchment level. This approach forces the model to evaluate environmental covariates and the latent spatial effect at a granularity that appropriately reflects the available data. Although this process is computationally intensive and prevents a global analysis, it illustrates the feasibility and utility of combining high-resolution spatial covariates with LGCPs at irregular spatial discretisations—something not commonly addressed in the literature (Liu and Vanhatalo, 2020; Diggle *et al.*, 2013).

The remainder of this section is organised as follows: Section 4.5.1 describes the dataset used in this analysis, Section 4.5.2 outlines the refined integration scheme, and Section 4.5.3 presents the modelling results.

4.5.1 Data

Japan, an island nation in the north-west Pacific Ocean, spans a total land area of approximately 377,000km^2, of which around 70% is mountainous (Yoshimatsu and Abe, 2006). Geologically, the country lies within the seismically active Ring of Fire—a tectonic belt characterised by subduction zones, deep oceanic trenches, and volcanic arcs. These features render Japan especially susceptible to natural hazards, including earthquakes, volcanic eruptions, and landslides. The socio-economic consequences of some of these events have been substantial, e.g. the Tohoku earthquake and tsunami of 2011, which triggered the Fukushima Daiichi nuclear disaster (Klein, 2023).

Landslides in Japan are commonly triggered by seismic activity, heavy rainfall, or a combination of both. Notable events include the 1984 Otaki earthquake, which triggered a high-velocity mudslide on Mount Ontake that killed 29 individuals. Similarly, Tropical Storm Talas in 2011 resulted in extensive infrastructural damage due to bedrock and soil failure from intense precipitation (Oku *et al.*, 2014). These climatic and geological conditions highlight the relevance of spatial modelling of landslide risk.

As noted earlier, this study focuses on a subset of Japan, specifically the central region. This choice was motivated by both computational and practical considerations. Japan is divided into approximately 300,000 catchments, and modelling the entire domain would be computationally prohibitive for the purposes of illustrating the *mess with the mesh* issue. Moreover, the central region exhibits a high concentration of landslide occurrences, making it a particularly relevant and informative area for focused analysis.

Figure 4.7 displays our area of study, the central region of Japan, and the zoom panels highlight the catchment scale. The catchment is an area of land which contributes water to a nominated point on the earth's surface (Bren, 2023), and a form that divides the study area in a way that preserves the geomorphological conditions relevant to land sliding processes (Amato *et al.*, 2019). The catchments are represented as irregular spatial polygons defined using `r.watershed` (GRASS Development Team, 2024) through Geographic Resources Analysis Support System (GRASS), a Geographic Information System (GIS).

The available explanatory covariates encompass a range of geographical and geological predictors. Notably, the dataset includes aggregated annual rainfall statistics, which serve as key indicators for rainfall-induced landslides, and mean surface velocity, a critical predictor for landslides triggered by seismic activity. Mean surface velocity reflects the maximum rate of ground movement, with higher values typically observed near earthquake epicentres and along fault lines. These elevated velocities

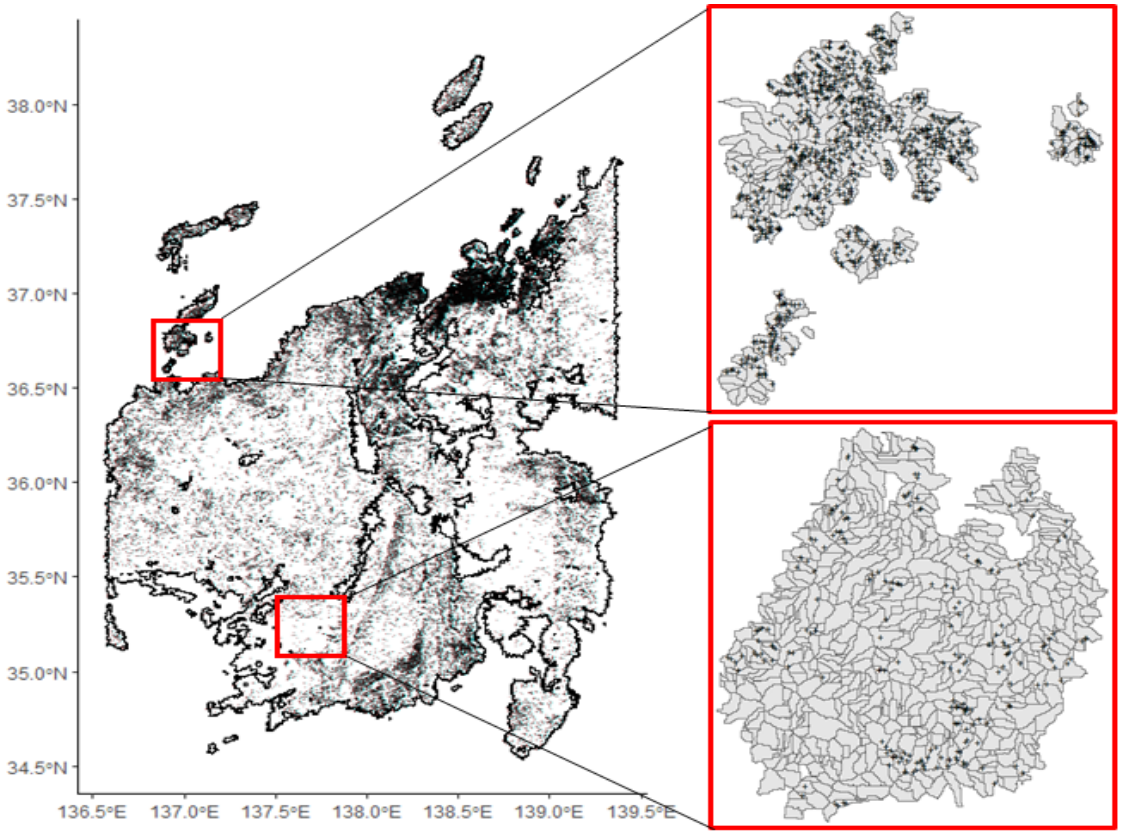

Figure 4.7: Main study subset of Japan with two example panels illustrating landslide point patterns from two of the 100 subsets used for model fitting.

are associated with an increased risk of structural damage and slope failure. Elevation is another important covariate, as mountainous regions inherently present greater susceptibility to landslides due to steeper terrain and gravitational forces. Figure 4.8 displays the spatial distribution of mean elevation and mean surface velocity across the catchments in central Japan.

The geographical and geological covariates included in the models were chosen using a stepwise forward selection procedure, allowing for both linear and non-linear effects (treated as fixed or random; see Chapter 1). At each step, model performance was assessed by comparing the fitted values for the mean count of landslides and for landslide size against their observed counterparts. This iterative approach helped determine whether a given covariate meaningfully improved the model. While information criteria such as the Deviance Information Criterion (DIC) and the Watanabe-Akaike Information Criterion (WAIC) are commonly used to guide covariate inclusion in standard regression models (Meyer, 2014; Gelman *et al.*, 2014), their applicability to point process models is limited due to the ambiguity surrounding sample size (Choiruddin *et al.*, 2021). In point process contexts, locations

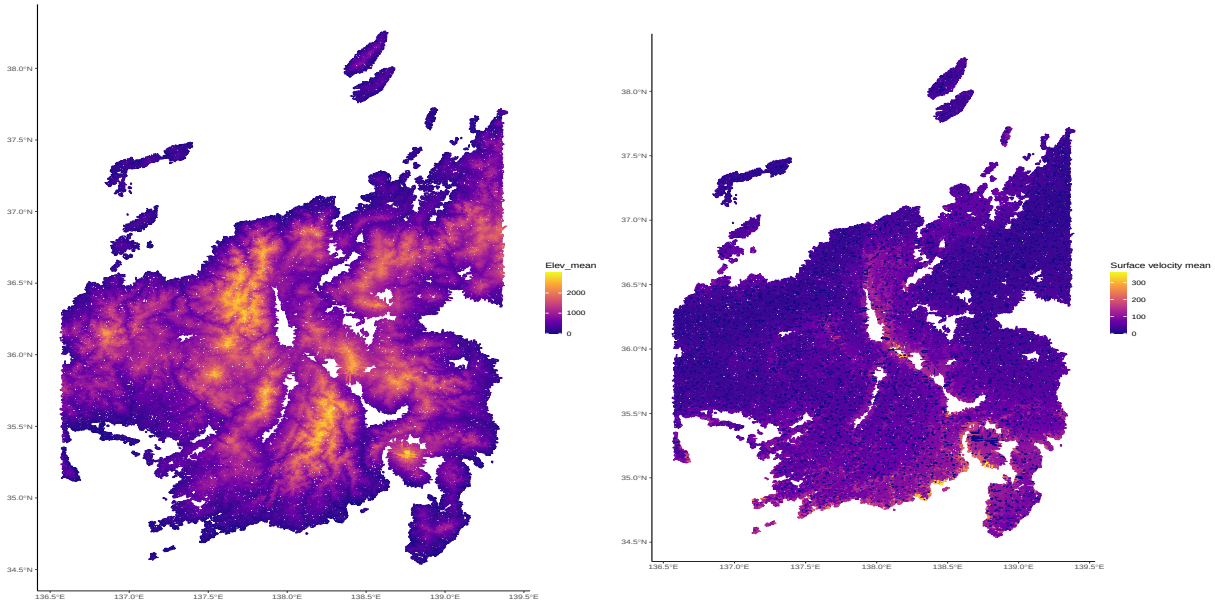

Figure 4.8: Spatial distribution of mean elevation (left) and mean surface velocity (right) across catchments in central Japan. These covariates are key predictors in landslide susceptibility modelling, with elevation capturing topographic variation and surface velocity reflecting seismic-related ground movement.

where points (events) have been observed as well as locations where no points have been observed carry informational value, and it is unclear how to quantify the amount of information that a point pattern reflects. In standard regression sample size is used to define these criteria, but this is not directly translatable into the context of point processes. The final set of selected covariates, listed in Table 4.1, are all included non-linearly using a Gaussian random walk of order two (RW2) with penalised complexity (PC) priors (Simpson *et al.*, 2017) set to be weakly informative such that the probability that the standard deviations ($1/\tau$) are greater than 1, is 0.01. We assume that the same spatial covariates influencing the occurrence of landslides also govern the distribution of landslide sizes.

4.5.2 Refined integration scheme

Figure 4.9 illustrates the design and progressive refinement of the manual integration scheme. The top-left panel shows the initial relatively coarse mesh, in which individual triangles do not adequately align with the underlying catchment boundaries, highlighting inadequate spatial resolution for accurate integration over such fine-scale units. The subsequent panels demonstrate a progressive refinement

Table 4.1: Selected geographical and geological covariates used in the landslide marked point process for central Japan.

Variable	**Acronym**	**Units**
Elevation mean	Elev_μ	m
Minimum curvature mean	MinCv_μ	1/m
Maximum curvature mean	MaxCv_μ	1/m
Aspect mean	Asp_μ	Degrees
Relief mean	Rlf_μ	m
Surface velocity mean	SVel_μ	m/s
Maximum rainfall mean	RnMax_μ	mm

of the integration scheme, achieved by incrementally increasing the number of integration points within each triangle. This refinement is implemented using the function `fm_subdivide(mesh, n)` as described in Section 4.4. This strategy is particularly recommended when applying the SPDE approach to spatial point process models involving high-resolution environmental covariates, as it improves spatial accuracy without necessitating a complete remeshing of the domain.

4.5.3 Results

Figures 4.10 to 4.12 display the estimated mean intensity surfaces for landslide occurrences, the corresponding estimated total counts, and QQ plots comparing observed and fitted log-landslide sizes for nine marked LGCPs fitted to subsets of our study area. These figures highlight the two key outputs of interest: the spatial intensity of events and the distribution of landslide sizes. As this case study is not the primary focus of the chapter, the results are presented as surface summaries, and detailed diagnostic assessments, such as residual analysis or formal validation, are omitted. Nonetheless, the models demonstrate good performance in capturing both the spatial pattern of landslide events and the distribution of log-transformed sizes. Overall, these results effectively demonstrate the successful integration of high-resolution spatial covariates into a marked point process model using the stochastic partial differential equation (SPDE; Lindgren *et al.*, 2011) - integrated nested Laplace approximation (INLA; Rue *et al.*, 2009) framework.

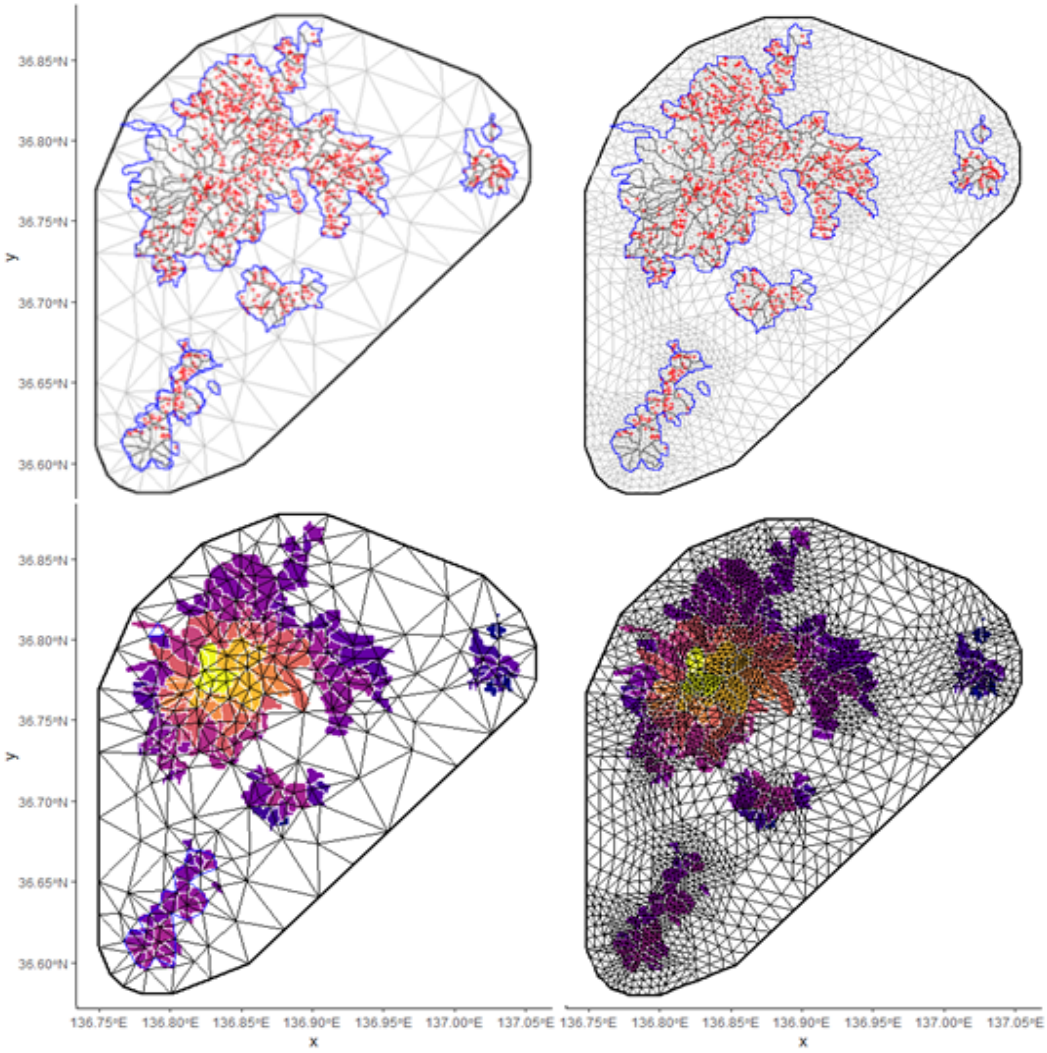

Figure 4.9: Illustration of increasing mesh refinement over a subset of the study area.

4.6 Discussion

In this chapter, we have identified and systematically examined the conditions under which the so-called *mess with the mesh* arises in spatial point process models using the SPDE-INLA framework. This issue stems from a mismatch between the spatial resolution of spatial covariates and the integration scheme induced by the triangulation mesh. Through two carefully designed simulation studies, we demonstrated how this mismatch can lead to substantial bias in the estimation of regression coefficients. The first simulation focused on shifts in the domain of integration, while the second explored the effects of mesh refinement. Both highlight how sensitive parameter estimates can be to the discretisation choices inherent in SPDE-based models.

To address these challenges in practical applications—such as landslide hazard modelling with high-resolution environmental data—we propose three key guidelines: (i) smooth the covariate surface when preserving fine-scale variability is not critical, to reduce sensitivity to mesh geometry; (ii) use the `fm_subdivide` function to manually refine the integration scheme without altering the original mesh, ensuring better alignment between mesh triangles and the resolution of the covariates; and (iii) incorporate basic model validation checks, such as comparing the expected and observed number of events across the domain, to assess the adequacy of the integration scheme. The first guideline requires the modeller to consider the purpose of their analysis. Fine-scale variability should be preserved when the covariate resolution carries important information for the interpretation or prediction of the process (e.g., when small-scale patterns drive hazard occurrence). Conversely, in exploratory analyses or when the precise resolution is not essential, covariates can be smoothed without compromising the study's objectives.

Our recommendations are illustrated in a case study from central Japan, where we successfully implemented a marked log-Gaussian Cox process using high-resolution covariate data. Overall, this chapter underscores the importance of aligning spatial discretisation strategies with the scale of environmental information and offers practical solutions to improve the robustness of point process modelling in applied spatial statistics.

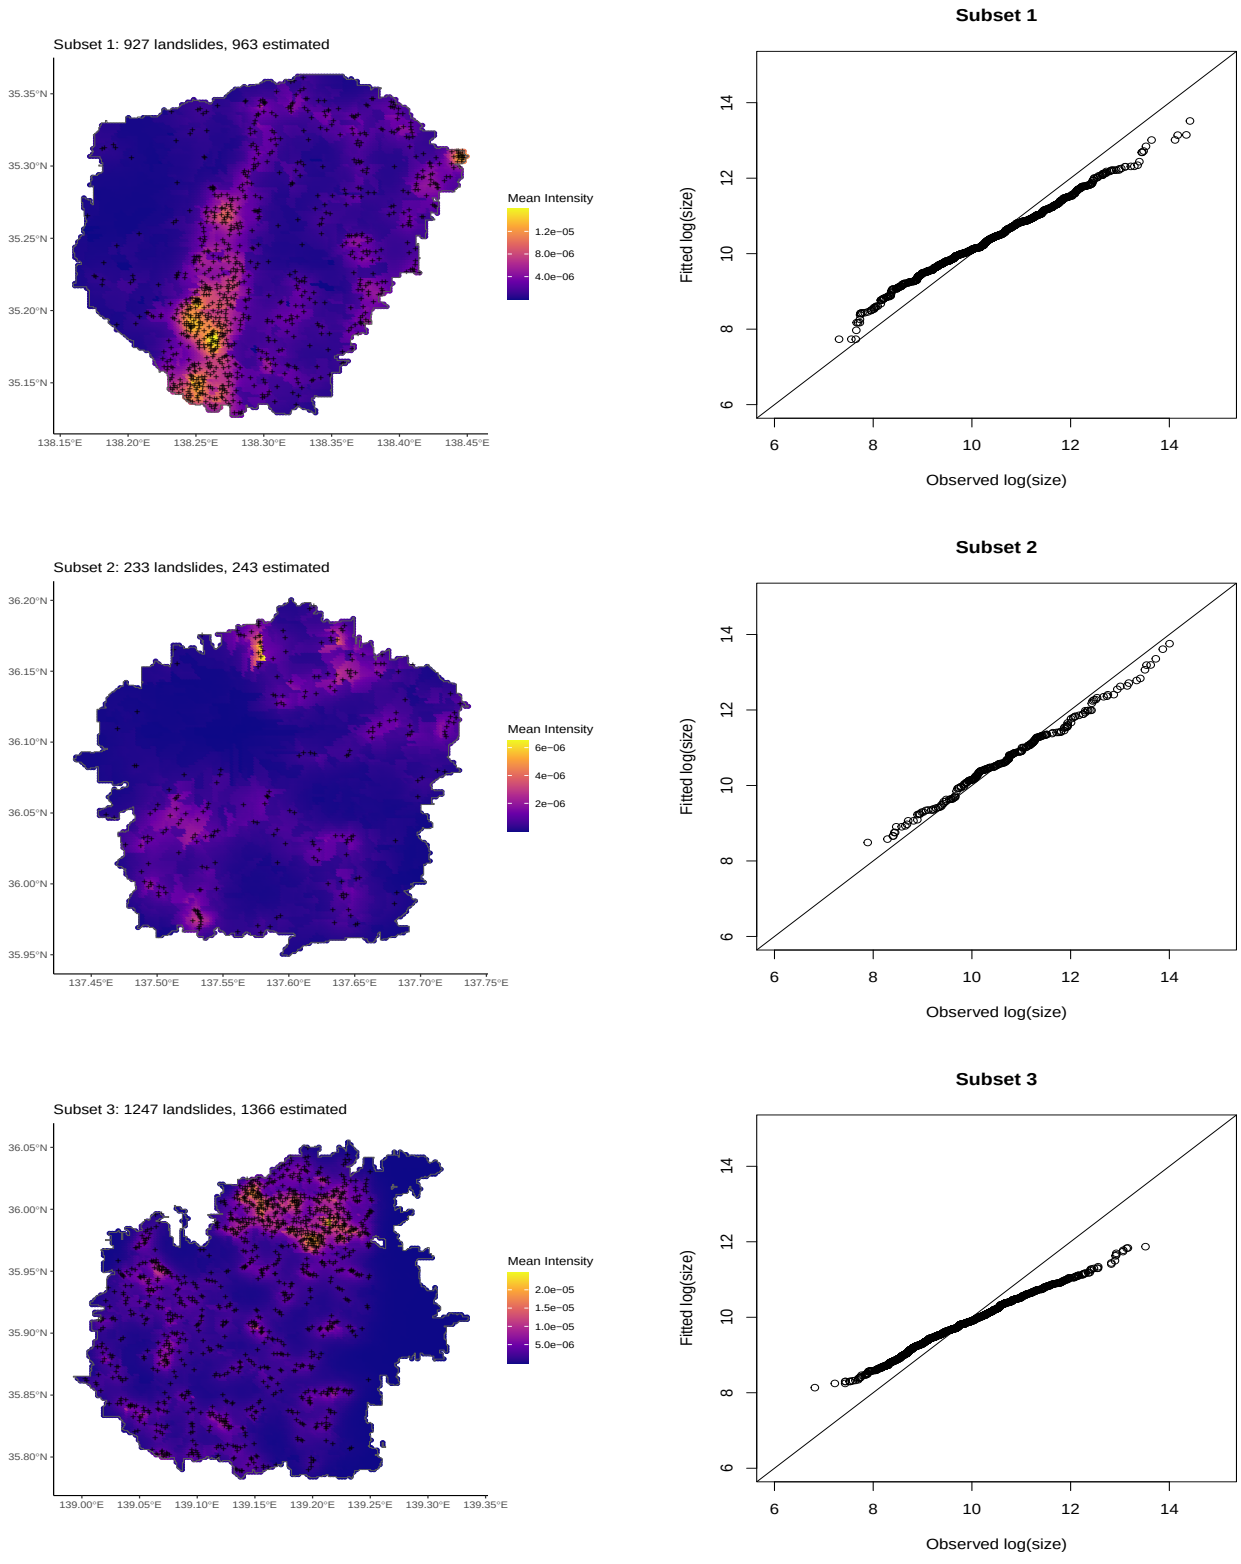

Figure 4.10: Estimated mean intensity fields with noted observed versus estimated landslide count (left column) and the QQ plot for the observed versus fitted landslide size distribution (right column) for subsets 1-3.

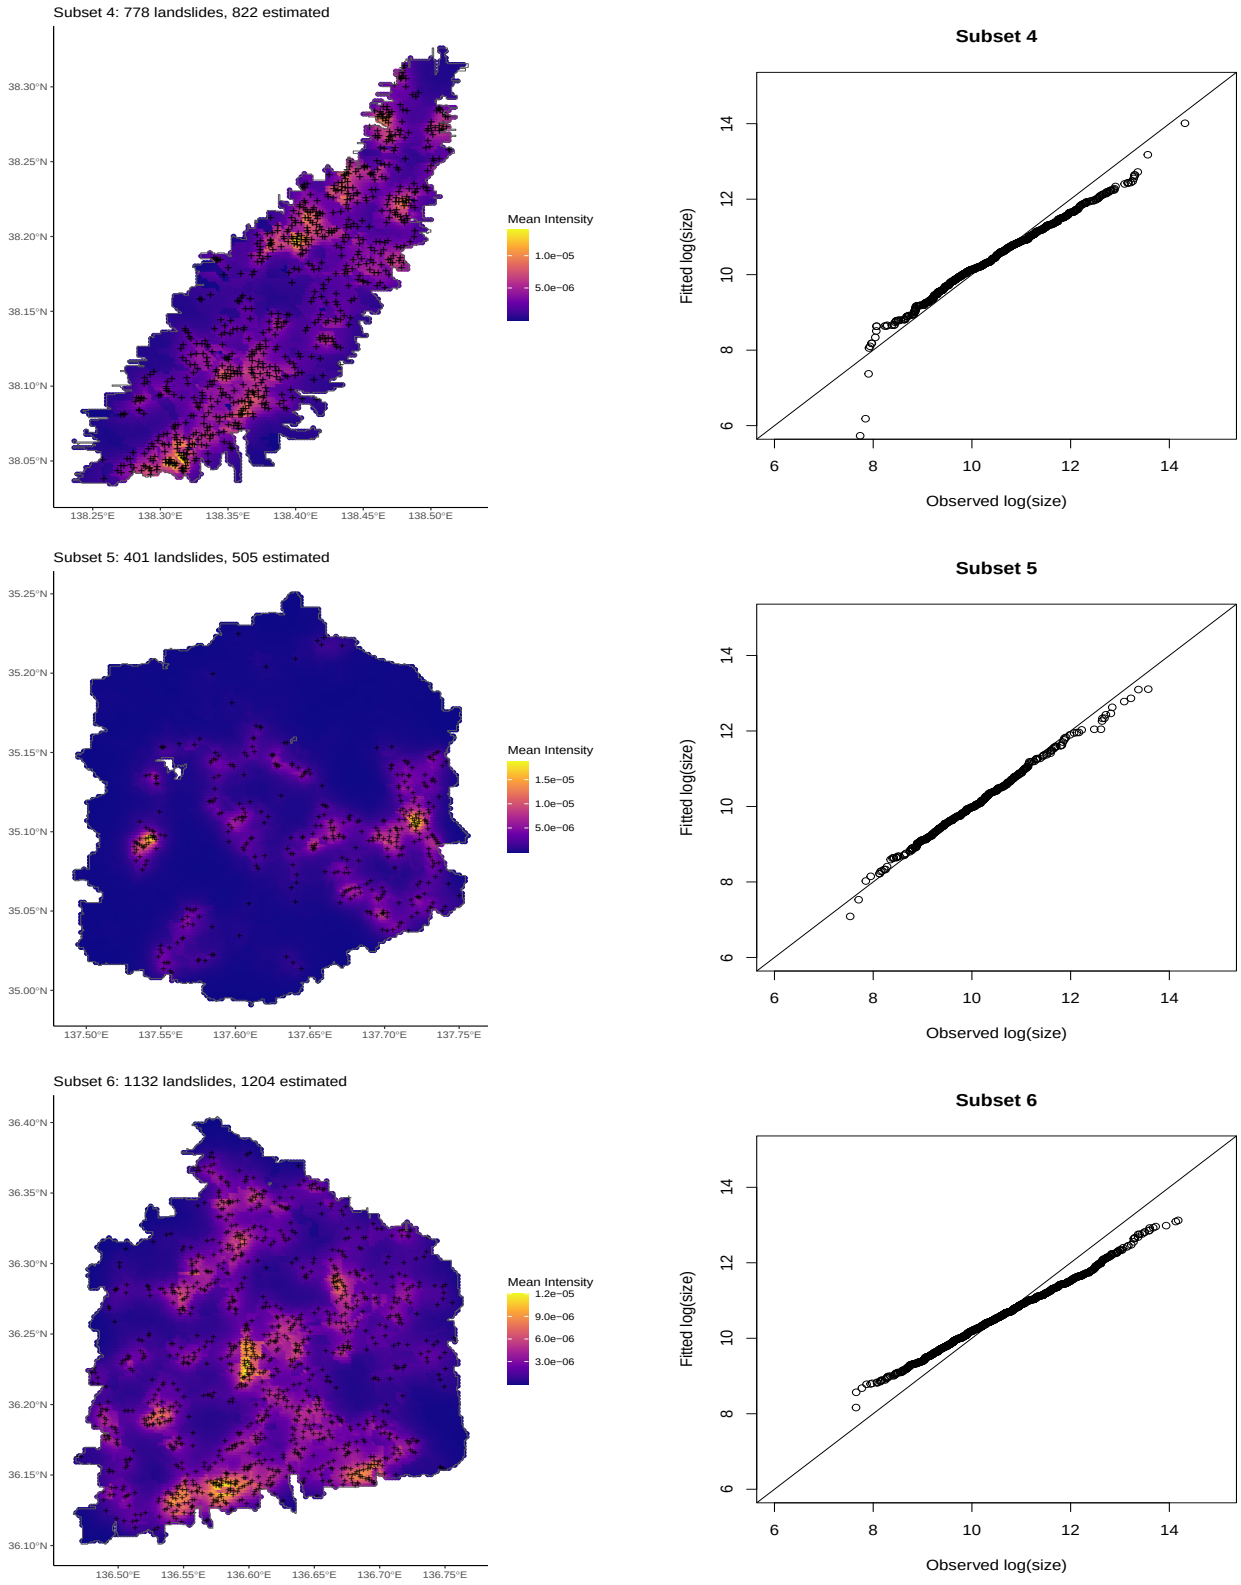

Figure 4.11: Estimated mean intensity fields with noted observed versus estimated landslide count (left column) and the QQplot for the observed versus fitted landslide size distribution (right column) for subsets 4-6.

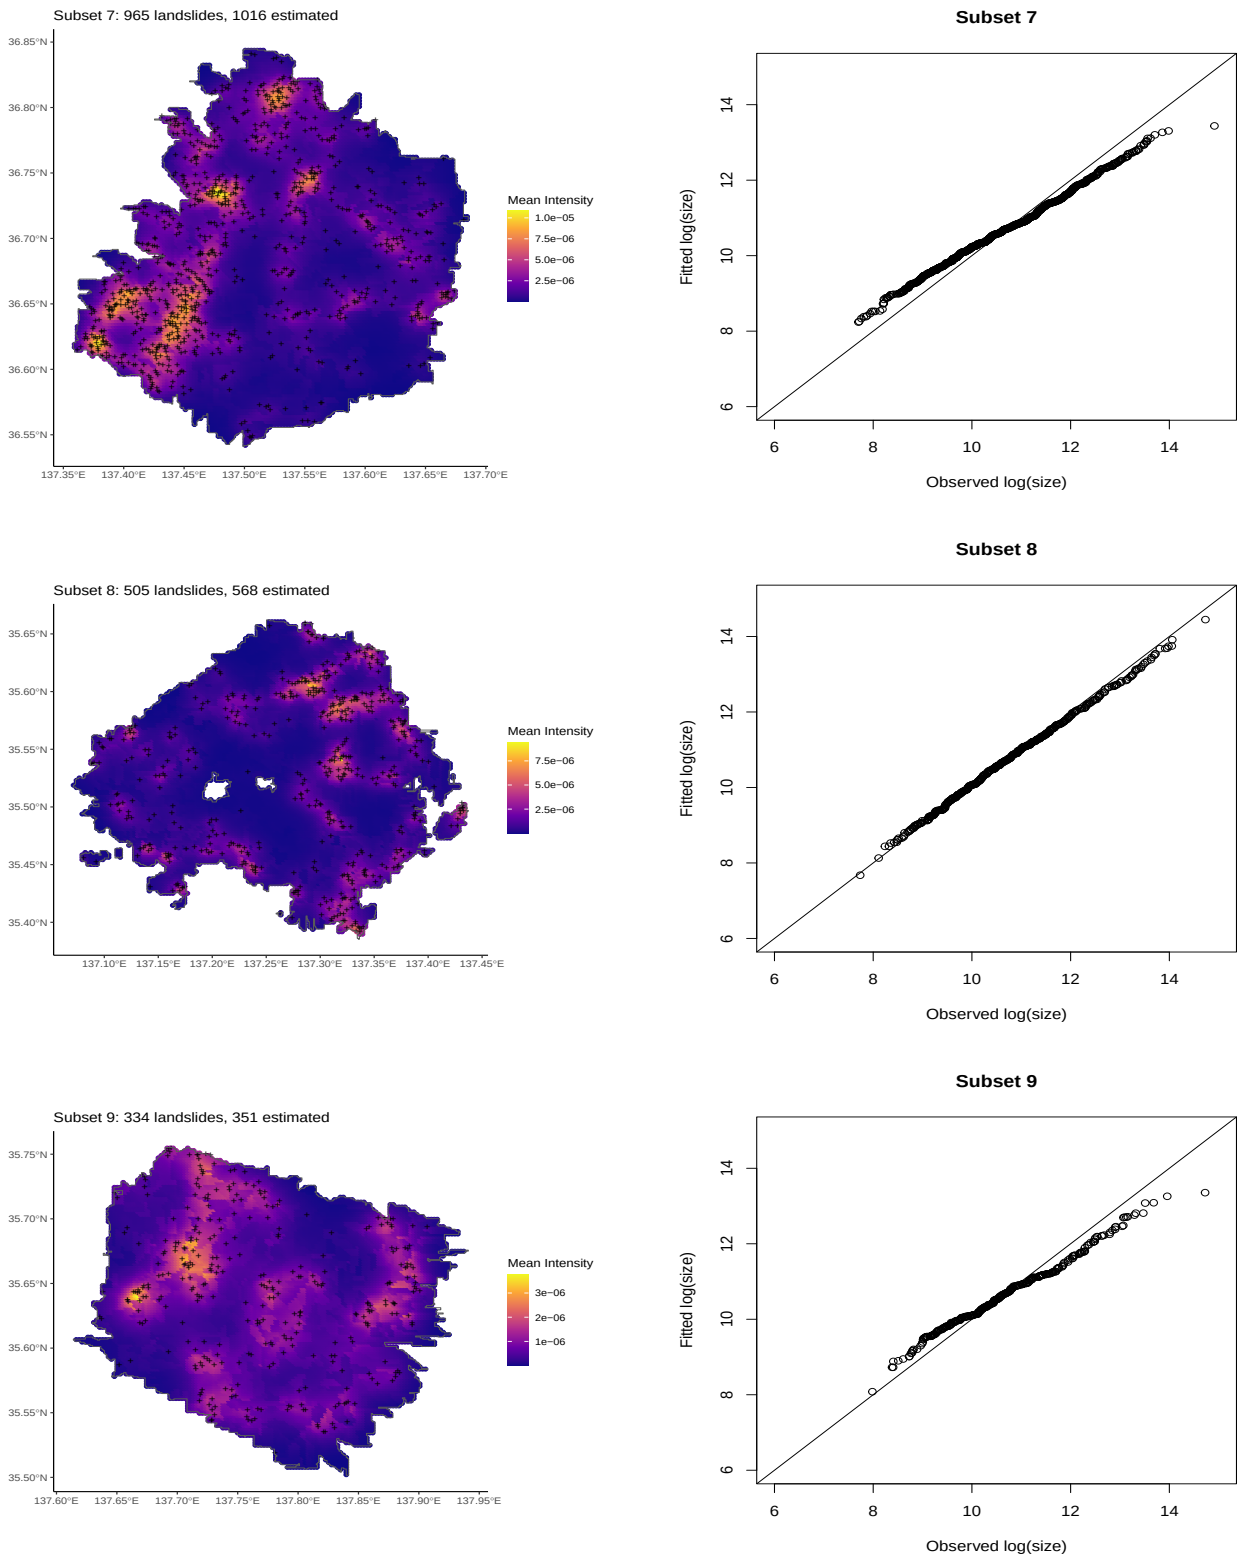

Figure 4.12: Estimated mean intensity fields with noted observed versus estimated landslide count (left column) and the QQplot for the observed versus fitted landslide size distribution (right column) for subsets 7-9.

Chapter 5

Spatio-temporal functional regression for landslide hazard modelling: the spatio-temporal SPDE as a smoothing penalty

This chapter extends previous landslide hazard analyses by incorporating the temporal dimension through spatio-temporal surface deformation data. It introduces a functional representation of daily precipitation as a predictor, enabling finer temporal resolution than traditional summary metrics. Building on Scalar-on-Function Regression (SoFR) within the mgcv framework, the chapter embeds a spatio-temporal Gaussian process in the linear predictor using a custom SPDE-derived penalty. This represents a novel integration of functional data analysis with separable spatio-temporal smoothing in a frequentist GAM setting.

5.1 Introduction

The work presented in this thesis so far has primarily addressed the question of *where* a landslide occurs and *how large* the associated landslide hazard is. In this chapter, we extend the analysis to include the temporal dimension. Specifically, we examine spatio-temporal measurements of ground surface deformation over a nineteen-month period in an area of the Sichuan province, China. Surface deformation refers to the extent of ground change, in this case sinking, and provides essential

information for landslide hazard assessment by revealing patterns of ground displacement that can indicate slope instability. We also revisit how precipitation is incorporated as a covariate in the model. Specifically, in Chapter 4, precipitation was represented using summary statistics, such as mean or maximum annual values, which are simplifications of a much richer, temporally detailed dataset. Here, we use daily precipitation data to construct a functional representation over time, capturing variation at a finer temporal resolution than the observed surface deformation. We aim to model precipitation as a continuous function across the days preceding and between surface deformation observations, to achieve a more nuanced characterisation of the deformation process, improving upon models based on aggregated precipitation metrics.

To implement a regression model with a functional predictor and a scalar response, we draw on the framework of Functional Data Analysis (FDA), focusing in particular on Scalar-on-Function Regression (SoFR). SoFR assumes that each observed function represents an individual, self-contained, and complete data point (Crainiceanu *et al.*, 2024). Rather than summarising a functional predictor using a single value, SoFR models assume a smooth relationship between the predictor and the response across the entire domain of the function. Estimation and inference are carried out using a Generalised Additive Model (GAM) with penalised regression splines, where the degree of smoothness is selected automatically via generalised cross-validation or restricted maximum likelihood (REML) (Wood, 2011). This approach is typically implemented in `R` using the `mgcv` package. As a result, we temporarily depart from our Bayesian inference framework in favour of a frequentist approach. This decision is motivated by the current limitations of INLA in handling functional covariates. While it is possible to include functional predictors in INLA by projecting them onto a set of basis functions and incorporating the resulting scores into the model—either linearly or non-linearly—this approach still involves a form of dimensionality reduction. Consequently, the functional predictor is not modelled as a continuous process, but rather through a finite set of coefficients, leading to a single estimated effect.

A key advantage of the Bayesian modelling framework in INLA is the seamless integration of the Stochastic Partial Differential Equation (SPDE) approach, which provides a computationally efficient way to model residual spatial dependence via Gaussian Random Fields (GRFs) (Lindgren *et al.*, 2011). This approach is particularly effective in capturing unexplained spatial correlation between observations while maintaining scalability through sparse precision matrices. In contrast, frequentist approaches offer several alternatives for modelling spatial effects. These include modelling spatial dependence using Gaussian Processes (GPs) estimated via maximum likelihood or restricted

maximum likelihood (REML; Haskard *et al.* 2007), employing penalised spline smoothers within the Generalised Additive Model (GAM) framework (Wood, 2017), and more recently, estimating SPDE models through likelihood-based methods on finite element meshes (Bolin and Kirchner, 2020).
The `mgcv` package offers a flexible framework for specifying complex regression models with smooth effects, including spatial and spatio-temporal components via tensor product smooths in a frequentist framework. However, while it includes low-rank approximations to Gaussian processes (e.g., via `bs = "gp"`), it does not allow for the explicit specification of general, user-defined covariance structures—particularly those that are non-stationary or non-separable in space and time. Spatio-temporal effects in `mgcv` are typically modelled using separable tensor product smooths (e.g., `te(x, y, t)`), which are analogous to separable spatio-temporal Gaussian processes. While this approach allows for interaction between space and time through a multiplicative structure, it is limited in its ability to capture more complex, non-separable dynamics where spatial and temporal dependencies interact in a non-multiplicative fashion. Moreover, the Gaussian process (GP) smooths available in `mgcv` (e.g., via `bs = "gp"`) are restricted to specific kernel families and rely on data-driven smoothing estimation, rather than the explicit specification of covariance structures or priors, as is typical in fully Bayesian GP models.

Given the spatio-temporal nature of our data, we adopt a space-time SPDE framework (Cameletti *et al.*, 2013), which allows for joint modelling of spatial and temporal dependence in a separable, structured way. We build on the work of Miller *et al.* (2020), who showed that spatial SPDEs can be expressed as smooth penalty matrices and implemented within the GAM framework using `mgcv`. In this chapter, we extend their approach to the spatio-temporal setting, demonstrating how the precision matrix derived from a spatio-temporal SPDE can be used as a penalty matrix within `mgcv` to construct a custom smoother that embeds a spatio-temporal Gaussian process in the linear predictor.

To our knowledge, this is the first implementation of a SoFR model in which the linear predictor includes a spatio-temporal Gaussian process specified through an SPDE-derived penalty matrix. This integration within the `mgcv` framework bridges functional data analysis with SPDE-based smoothing in a computationally efficient and modular way. It introduces a novel class of models relevant in, e.g. complex environmental and geostatistical applications, where both functional predictors and a spatio-temporal structure play a central role.

The remainder of this chapter is organised as follows. Section 5.2 presents the available data and outlines the Interferometric Synthetic Aperture Radar (InSAR) technique to measure ground surface

deformation. Section 5.3 introduces the functional representation of precipitation, the proposed space-time SPDE smoother, and details its implementation within the `mgcv` framework. Section 5.5 reports the model results, and Section 5.6 concludes with a discussion of the key findings and future directions.

5.2 Data

Interferometric Synthetic Aperture Radar (InSAR) is a remote sensing technique used to measure ground surface deformation with high spatial resolution and accuracy. It operates by comparing radar signals during multiple passes over the same area at different times. The reflected signals contain phase information that corresponds to the distance between the satellite and the ground. By analysing the phase difference between two or more SAR images taken at different times, InSAR can detect minute changes in the Earth's surface—often on the order of millimetres. We have 163 images (He *et al.*, 2025). This inventory features ground surface deformation extracted through InSAR, across a region in Sichuan, China on every 12^{th} day between April 2017 and November 2018. Figure 5.1 illustrates a sample of three temporal observations of surface deformation across the space. The

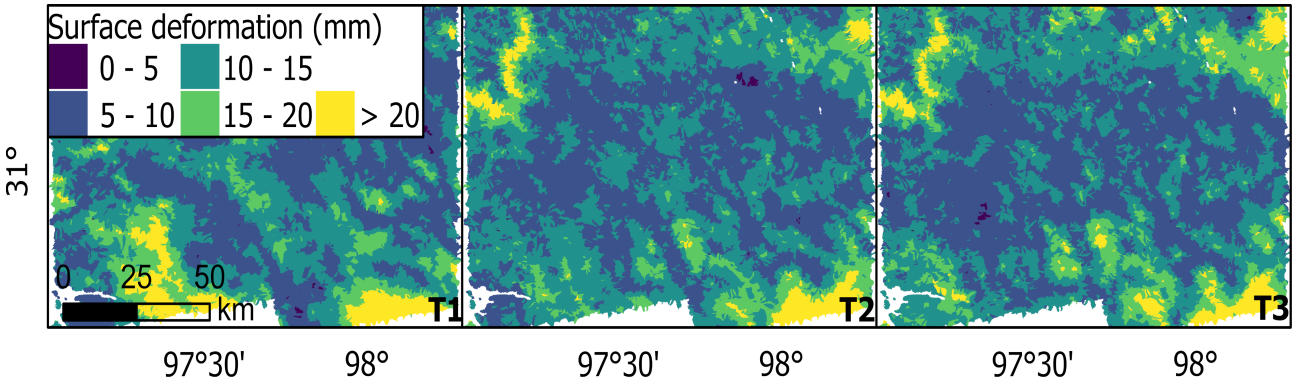

Figure 5.1: Surface deformation measurements across the region in Sichuan, China for three time points. T1) day zero, T2) day 12 and T3) day 24.

spatial information of the study area was aggregated through a Slope Unit (SU, hereafter) partition, calculated through `r.slopeunits` (Alvioli *et al.*, 2016). The SU encompasses the geographic space between streams and ridges (Amato *et al.*, 2019), providing a unit for interpretation purposes, and we have a total of 23,140 SUs. Figure 5.2 provides a visual representation of these SUs through geographical and geological predictors. From the available covariates, we selected a subset based on expert knowledge rather than relying on automated stepwise procedures. This choice reflects the primary aim of the analysis: to explore the integration of a functional covariate and a spatio-temporal

SPDE smoother. Table 5.1 summarises the chosen covariates and specifies whether each was modelled with a linear or smooth effect. These covariates include both geographical and geological features that characterise the morphological complexity of the Chinese landscape. Geographical variables are treated as continuous, while geological characteristics are incorporated as categorical factors. Figure 5.3 displays the functional representation of precipitation over time for each SU.

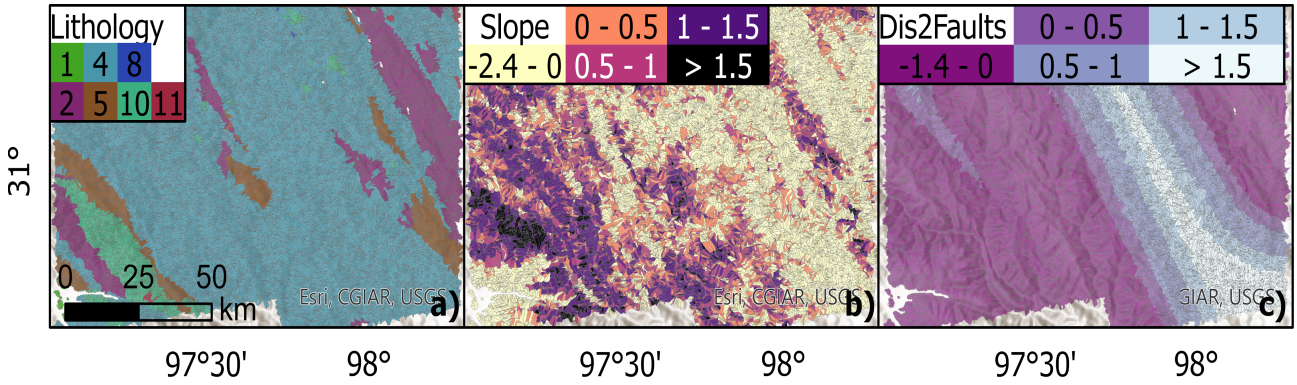

Figure 5.2: Slope Unit discretisation of the region in Sichuan, China with geographical descriptors, a) lithology classes, b) average standardised slope angle, and c) average standardised distance to the nearest fault-line.

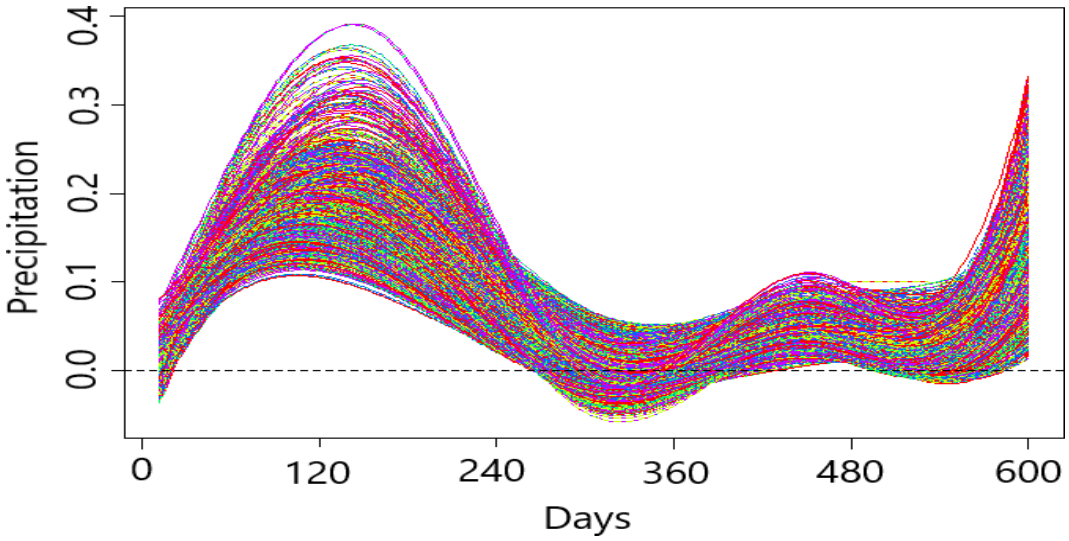

Figure 5.3: Functional precipitation measurements for each Slope Unit.

Table 5.1: Selected covariates for inclusion in the model and their associated form of entry.

Variable	**Units**	**Acronym**	**Effect**
Precipitation	mm	Precip	functional
Slope mean	Degrees	Slo_μ	smooth
Mean distance to nearest fault-line	m	Dist2Faults_μ	smooth
Mean distance to nearest river	m	Dis2River_μ	smooth
Lithology class	n/a	Lith	linear - factor
Profile curvature mean	1/m	ProfC_μ	linear
Planar curvature mean	1/m	PlanC_μ	linear

5.3 Methodology

Let $Y(\mathbf{s}, \mathbf{t})$ be the surface deformation in $\mathbf{s} \in \mathcal{S} \subset \mathbb{R}^2$, i.e. the set of 23,140 slope units, and $\mathbf{t} \in \mathcal{T} = \{0, 12, 24, ..., 588\}$, the discrete time points at 12-day intervals. We assume

$$\log Y(\mathbf{s}, \mathbf{t}) \sim \mathcal{N}(\mu(\mathbf{s}, \mathbf{t}), \sigma^2),$$

where $\mu(\mathbf{s}, \mathbf{t})$ is the mean surface deformation and σ^2 is the variance. The linear predictor is given by

$$\eta(\mathbf{s}, \mathbf{t}) = \alpha + \sum_{m=1}^{M} \beta_m x_m(\mathbf{s}, \mathbf{t}) + \sum_{k=1}^{K} f_k(w_k(\mathbf{s}, \mathbf{t})) + z(\mathbf{s}, \mathbf{t}) + u(\mathbf{s}, \mathbf{t}), \tag{5.1}$$

and is linked to the mean using the default identity link: $\mu_{\mathbf{s}}(\mathbf{t}) = \eta_{\mathbf{s}}(\mathbf{t})$. In 5.1, α is an intercept and $(x_1(\mathbf{s}), \ldots, x_M(\mathbf{s}))^\top$ are the linear effect covariates described in Table 5.1 with fixed coefficients $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_M)^\top$. The functions $\boldsymbol{f} = \{f_k(\cdot), \ldots, f_K(\cdot)\}$ are smooth effects defined in terms of the smooth effect covariates $(w_1(\mathbf{s}), \ldots, w_K(\mathbf{s}))^\top$ in Table 5.1, the term $z(\mathbf{s}, \mathbf{t})$ is the precipitation (see Section 5.3.1), and $u(\mathbf{s}, \mathbf{t})$ is the random spatio-temporal effect (see Section 5.3.2).

5.3.1 Precipitation term

We now examine the precipitation covariate in more detail, as it enters the model as a functional predictor. We observe daily precipitation, denoted $X_{\mathbf{s}}(p)$, where $\mathbf{s} \in \mathcal{S}$ is a slope unit and $p \in \mathcal{P} = \{0, 1, 2, \ldots, 588\}$ indexes daily time points. For each deformation observation at time $t \in \mathcal{T}$, we define the local precipitation domain as

$$P_t = \{p \in \mathcal{P} : t - \delta < p \leq t\},$$

where δ specifies the number of days used to construct the functional precipitation predictor (e.g., $\delta = 12$ for a 12-day precipitation history). Thus, for each $t \in \mathcal{T}$, precipitation is treated as a function defined on the interval P_t, representing a temporally resolved history of precipitation leading up to the deformation observation. This construction allows us to incorporate precipitation in the model using a Functional Generalised Additive Model (FGAM) term:

$$z(\mathbf{s}, \mathbf{t}) = \int_{P_t} F(X_{\mathbf{s}}(p)), p)dp = \int_{t-\delta}^{t} F(X_{\mathbf{s}}(p)), p)dp,$$

where $F(x, p)$ is a smooth bivariate function capturing the interaction between the magnitude of precipitation and its timing, and $X_{\mathbf{s}}(p)$ is the observed functional predictor. Following McLean *et al.* (2014), we approximate $F(x, p)$ using a tensor-product of B-spline bases:

$$F(x, p) = \sum_{j=1}^{K_x} \sum_{k=1}^{K_p} \alpha_{j,k} \mathbf{B}_j^X(x) \mathbf{B}_k^P(p) \tag{5.2}$$

where $\mathbf{B}_j^X(x)$ and $\mathbf{B}_k^P(p)$ are spline basis functions defined over the domains of precipitation magnitude and time, respectively. Substituting into the integral yields:

$$z(\mathbf{s}, \mathbf{t}) = \sum_{j=1}^{K_x} \sum_{k=1}^{K_p} \alpha_{j,k} \int_P \mathbf{B}_j^X(X_{\mathbf{s}}(p)) \mathbf{B}_k^P(p)dp \tag{5.3}$$

Since precipitation is recorded at discrete time points, we approximate the integral numerically:

$$z(\mathbf{s}, \mathbf{t}) \approx \sum_{j=1}^{K_x} \sum_{k=1}^{K_p} \alpha_{j,k} \sum_{i=1}^{n} \mathbf{B}_j^X(X_{\mathbf{s}}(p_i)) \mathbf{B}_k^P(p_i) v_i, \tag{5.4}$$

where $\{p_i\}_{i=1}^n \subseteq P_t$ are the daily observation points within the interval P_t, and v_i are numerical integration weights (Scheipl *et al.*, 2015). This formulation corresponds to a tensor-product smooth over the space of functional covariates and their temporal support, where $\alpha_{j,k}$ are the coefficients to be estimated.

5.3.2 Space-time SPDE penalty smoother

This section outlines how the SPDE approach translates into a smoothing penalty, beginning with the spatial case and then extending to the spatio-temporal setting relevant to our model for surface deformation in China.

An SPDE relates differential operators D (e.g., first or second derivatives, gradient, or Laplacian operators) to a stochastic process, typically white noise. Specifically, it takes the form $Df = \epsilon$, where ϵ is Gaussian white noise. Lindgren *et al.* (2011) proposed a computationally efficient approximation

to solutions of such SPDEs by leveraging the Finite Element Method (FEM; Brenner 2008). FEM discretises the spatial domain $\mathcal{S}$ using a triangulated mesh, whose vertices define a finite-dimensional basis for representing the spatial field. Let $\psi_k(\mathbf{s})$ denote a set of piecewise linear basis functions associated with the K vertices of the mesh. The spatial field $u(\mathbf{s})$, which is a realisation of a Gaussian Random Field (GRF), is then approximated as

$$u(\mathbf{s}) \approx \sum_{k=1}^{K} \psi_k(\mathbf{s}) w_k,$$

where the weights $w_k \sim \mathcal{N}(0, Q^{-1})$ are Gaussian random variables, and Q is a precision matrix encoding the spatial dependence structure. The values of the field at mesh vertices are determined by the weights, while values in the interior of the mesh elements are obtained via linear interpolation of the basis functions.

To generate a GRF with Matérn covariance (Matérn, 1986a), the SPDE formulation defines the differential operator as

$$D = \tau(\kappa^2 - \Delta)^{\alpha/2}, \tag{5.5}$$

where Δ is the Laplacian, τ is a scaling parameter, κ controls the spatial range, and $\alpha = \nu - d/2$, with ν being the smoothness parameter of the Matérn class and d the spatial dimension (here, $d = 2$). For computational convenience, we assume $\alpha = 2$, which yields a linear differential operator (Zhang, 2004). The FEM approach yields a sparse precision matrix Q that approximates the inverse of the Matérn covariance matrix. As shown by Miller *et al.* (2020), this precision matrix is equivalent to the penalty matrix derived from a smoothing spline formulation, thereby linking the SPDE representation to penalised smoothing in GAMs.

Smoothing penalties in GAMs are often formulated through the basis-penalty approach (Ramsay *et al.*, 1997), where the roughness of a function is penalised to avoid overfitting. The general form of such a penalty is

$$J(\boldsymbol{\beta}, \lambda) = \lambda \int \{Df\}^2 d\mathbf{s} = \lambda \langle Df, Df \rangle,$$

where D is a differential operator, β are the model parameters, and λ is a smoothing parameter. The role of the penalty is to favour smooth solutions by penalising high curvature or irregular behaviour in f. Hence, an SPDE of the form $Df = \epsilon$ implies a Matérn penalty of the form:

$$\langle Df, Df \rangle = \tau \int (\kappa^2 f - \Delta f)^2 d\mathbf{s},$$

as shown by Miller *et al.* (2020). We now generalise this to the spatio-temporal case.

Following Lindgren *et al.* (2022), the SPDE governing the spatio-temporal process is:

$$\left(a + \frac{\delta}{\delta t}\right)(\kappa^2 - \Delta)^{\alpha/2} f = \frac{\sqrt{(2a)}}{\tau}\mathcal{W}f,$$

where $\mathcal{W}$ is spatio-temporal white noise, and $f(\mathbf{s}, t)$ is now a function of both space and time. Defining $L = \kappa^2 - \Delta$, the differential operator becomes

$$D_{st} = \tau\left(a + \frac{\partial}{\partial t}\right)L. \tag{5.6}$$

Under this operator, the corresponding spatio-temporal penalty in strong form is

$$\langle D_{st}f, D_{st}f\rangle = \tau\int_{\mathcal{T}}\int_{\mathcal{S}}\left(a + \frac{\partial}{\partial t}\right)^2 (Lf)^2\, d\mathbf{s}, d\mathbf{t}. \tag{5.7}$$

To evaluate this, we first expand the operator:

$$\left(a + \frac{\partial}{\partial t}\right)^2 = a^2 + 2a\frac{\partial}{\partial t} + \left(\frac{\partial}{\partial t}\right)^2. \tag{5.8}$$

Separately, using Green's identity, Miller *et al.* (2020) show that

$$\langle Lf, Lf\rangle = \kappa^4\langle f, f\rangle + 2\kappa^2\langle \nabla f, \nabla f\rangle + \langle \Delta f, \Delta f\rangle. \tag{5.9}$$

Substituting (5.8) and (5.9) into (5.7) gives:

$$\begin{aligned}\langle D_{st}f, D_{st}f\rangle = \tau\Big\{a^2\left(\kappa^4\langle f, f\rangle + 2\kappa^2\langle \nabla f, \nabla f\rangle + \langle \Delta f, \Delta f\rangle\right)\\ + 2a\left\langle \frac{\partial f}{\partial t}, (Lf)^2\right\rangle\\ + \left\langle \left(\frac{\partial f}{\partial t}\right)^2, (Lf)^2\right\rangle\Big\}.\end{aligned}$$

This decomposition, illustrated in Figure 5.4 consists of three components:

- Spatial smoothing (first line),
- Spatio-temporal interaction (second line),
- Non-linear space-time coupling (third line).

Because this formulation is given in strong form, it requires sufficient smoothness of f, specifically that $f \in C^2(\mathcal{S}) \cap C^1([0, \mathcal{T}])$, i.e., twice differentiable in space and once in time (Sullivan, 2020). In practice, this regularity may be too restrictive for real data. As a result, the SPDE approach is typically implemented in weak form (variational formulation), in which we multiply by a test function $v(\mathbf{s}, t)$, integrate over space and time, and (if needed) apply integration by parts (Bakka *et al.*, 2018). This transfers derivatives from f to v, and only requires that f be weakly differentiable, relaxing the smoothness assumptions. In the following, we explore this variational formulation.

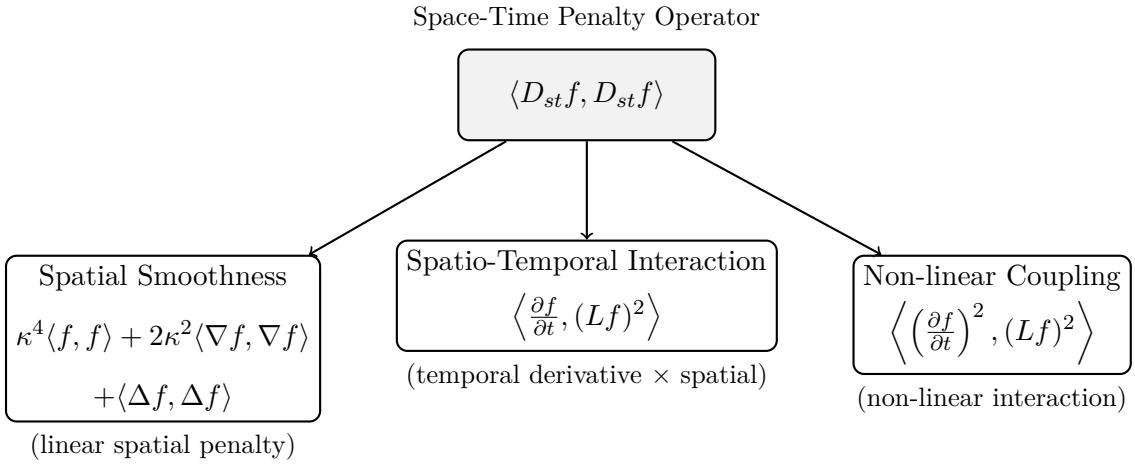

Figure 5.4: Decomposition of the space-time penalty operator $\langle D_{st}f, D_{st}f\rangle$ into spatial, spatio-temporal interaction, and non-linear coupling components.

Variational Formulation of $\langle D_{st}f, D_{st}f\rangle$

To derive the weak form of the spatio-temporal penalty, we multiply the strong form operator by a test function $v(s,t)$ and integrate over space and time. The resulting variational expression is:

$$\int_{\mathcal{T}}\int_{\mathcal{S}} \left(a + \frac{\partial f}{\partial t}\right)(Lf)\cdot v\, d\mathbf{s}\, dt,$$

where $L = \kappa^2 - \Delta$ and v is a test function defined over space and time. In the following, we consider the weak formulation for each dimension separately.

Spatial part. We begin with the spatial inner product

$$\int_{\mathcal{S}}(Lf)\cdot v\, d\mathbf{s} = \kappa^2\int_{\mathcal{S}} f\cdot v\, d\mathbf{s} - \int_{\mathcal{S}}\Delta f\cdot v\, d\mathbf{s}.$$

Applying Green's identity under homogeneous Neumann boundary conditions, the second term becomes:

$$-\int_{\mathcal{S}}\Delta f\cdot v\, d\mathbf{s} = \int_{\mathcal{S}}\nabla f\cdot\nabla v\, d\mathbf{s},$$

so the weak form of the spatial operator is:

$$\int_{\mathcal{S}}(Lf)\cdot v\, d\mathbf{s} = \kappa^2\int_{\mathcal{S}} f\cdot v\, d\mathbf{s} + \int_{\mathcal{S}}\nabla f\cdot\nabla v\, d\mathbf{s}.$$

In the finite element setting, we represent f and v using basis functions ψ_i, ψ_j as:

$$f = \sum_i f_i\psi_i, \quad v = \sum_j v_j\psi_j.$$

This yields the FEM system matrix:

$$K = \kappa^2 C + G_1,$$

where

- $C_{ij} = \langle \psi_i, \psi_j \rangle = \int_{\mathcal{S}} \psi_i \psi_j \, d\mathbf{s}$ (mass matrix),
- $(G_1)_{ij} = \langle \nabla \psi_i, \nabla \psi_j \rangle = \int_{\mathcal{S}} \nabla \psi_i \cdot \nabla \psi_j \, d\mathbf{s}$ (stiffness matrix).

For $\alpha = 2$, the second-order penalty is:

$$Q_s = KC^{-1}K = (\kappa^2 C + G_1)C^{-1}(\kappa^2 C + G_1) = \kappa^4 C + 2\kappa^2 G_1 + G_2,$$

where $G_2 = G_1 C^{-1} G_1$ approximates the inner product $\langle \Delta\psi_i, \Delta\psi_j \rangle$.

Temporal part. We now consider the temporal operator:

$$\int_{\mathcal{T}} \left(a + \frac{\partial f}{\partial t} \right) (Lf) \cdot v \, dt.$$

Assuming the spatial part $Lf \cdot v$ is already handled via the FEM matrices, this becomes a standard inner product in time. In practice, we discretise time with a uniform grid $t_1, \ldots, t_T$ and time step Δt, and model the temporal evolution of $f(t)$ using a first-order autoregressive process:

$$f(t_k) = \phi f(t_{k-1}) + \varepsilon(t_k), \quad \varepsilon(t_k) \sim \mathcal{N}(0, Q_s^{-1}), \tag{5.10}$$

with $\phi = 1 - a\Delta t$. The innovation term $\varepsilon(t_k)$ inherits the spatial precision structure. The differential operator $\left(a + \frac{\partial}{\partial t}\right)$ corresponds to a backward finite difference:

$$\left(a + \frac{\partial}{\partial t} \right) f(t_k) \approx \left(a + \frac{1}{\Delta t} \right) f(t_k) - \frac{1}{\Delta t} f(t_{k-1}).$$

FEM Implementation

We now assemble the full spatio-temporal precision matrix. Let:

- $Q_s = \kappa^4 C + 2\kappa^2 G_1 + G_2$: spatial precision matrix,
- $Q_t = \frac{1}{\Delta t^2} + a^2 I$: temporal precision matrix (based on AR(1) prior structure).

Then the full precision matrix is given by a Kronecker product:

$$Q = \tau Q_t \otimes Q_s = \tau \left(\frac{1}{\Delta t^2} + a^2 I \right) \otimes \left(\kappa^4 C + 2\kappa^2 G_1 + G_2 \right). \tag{5.11}$$

This formulation enables scalable computation for spatio-temporal models using sparse precision matrices and separable structures, while still allowing for interaction between spatial and temporal components through the joint penalty formulation.

5.4 Fitting the Matérn space-time SPDE and functional predictor in `mgcv`

The `mgcv` package in R allows the definition of custom smoothers through user-supplied `smooth.construct` functions. These return the model matrix $\mathbf{X}$ (evaluations of basis functions), a list of penalty matrices $\mathbf{S}$, and additional components required for smoothing parameter estimation.

In this section, we describe how to implement the space-time SPDE model with a functional covariate using a constructor `smooth.construct.spde.st.smooth.spec`. Let the response vector be $\{y_i\}_{i=1}^n$ and assume an $n \times n_c$ matrix of scalar covariates $\mathbf{X}_c$. The functional covariate is denoted $X_{\mathbf{s}}(p)$, observed at a grid of daily time points p.

Step-by-step Implementation

1. **Construct the spatial mesh:** Define a triangulation over the spatial domain using `INLA::inla.mesh.2d`, which determines the spatial basis functions $\{\psi_k(\mathbf{s})\}$.

2. **Define the temporal precision matrix:** Specify the AR(1) temporal dependence via its precision matrix $\mathbf{Q}_t$, with entries derived from a backward difference approximation of $(a + \partial/\partial t)$.

3. **Compute FEM matrices:** Use `INLA::inla.mesh.fem` to compute the standard FEM matrices:

   - Mass matrix $C = \int \psi_i \psi_j \, d\mathbf{s}$,
   - Stiffness matrix $G_1 = \int \nabla\psi_i \cdot \nabla\psi_j \, d\mathbf{s}$,
   - Higher-order matrix $G_2 = G_1 C^{-1} G_1$.

4. **Assemble the space-time penalty:** Construct the space-time SPDE penalty matrix using the Kronecker product (see (5.11)):

$$\mathbf{S}_{\text{spde}} = \tau \, \mathbf{Q}_t \otimes (\kappa^4 C + 2\kappa^2 G_1 + G_2).$$

5. **Build the SPDE model matrix:** Use `INLA::inla.spde.make.A(mesh, loc, group)` to obtain the projection matrix $\mathbf{A}$ that maps the latent field to observation locations. The linear predictor associated with the SPDE component is then given by $\mathbf{Aw}$, where $\mathbf{w}$ contains the latent weights.

6. **Represent the functional covariate:** Represent the daily precipitation process $X_{\mathbf{s}}(p)$ over its local support window $P_t = \{p \in \mathcal{P} : t - \delta < p \leq t\}$ using (5.2)-(5.4).

7. **Combine all model terms:** The full model matrix is then:

$$\mathbf{X} = [\mathbf{X}_c \;\; \mathbf{A} \;\; \mathbf{Z}],$$

where $\mathbf{X}_c$ is the matrix of scalar covariates, $\mathbf{A}$ is the SPDE model matrix (step 5.), and $\mathbf{Z}$ corresponds to the design matrix for the functional covariate obtained from the discretised FGAM integral. The full list of penalty matrices includes both:

- $\mathbf{S}_{\text{spde}}$: space-time SPDE penalty,
- $\mathbf{S}_{\text{fgam}}$: penalty matrix associated with smoothness of $F(x, p)$.

Both are passed to `mgcv` for joint estimation of smoothing parameters.

5.5 Results

This section presents the results of our regression model in terms of the statistical findings, and particular attention is given to the precipitation effect and the spatio-temporal SPDE smoother. Figure 5.5 shows the fitted means and the corresponding standard error intervals of the linear effects. The variable ProfCμ shows a positive association with surface deformation, though the large standard error suggests considerable uncertainty. In contrast, PlanCμ appears to have a negligible effect. All lithology classes are associated with positive estimated effects, with L2 and L4 standing out as the most influential. However, the standard error intervals for all lithology classes are wide and span both negative and positive values, indicating a lack of statistical significance at conventional levels. It is important to note that non-significance does not imply a true absence of effect, but rather reflects uncertainty in estimation. Figure 5.6 displays the fitted means and the corresponding standard error intervals of the smooth effects. It is essential to note that the variables were standardised before inclusion, which is reflected in the x-axis labels. We can see the highly non-linear influence of most of these covariates on surface deformation. For example, steeper slope angles are associated with increased deformation. While this may seem counterintuitive, it is consistent with findings that low-angle slopes may be more susceptible to vertical deformation during early failure (Sun *et al.*, 2022). The effect of Dist2Faultsμ is extremely variable and exhibits marked non-linearity.

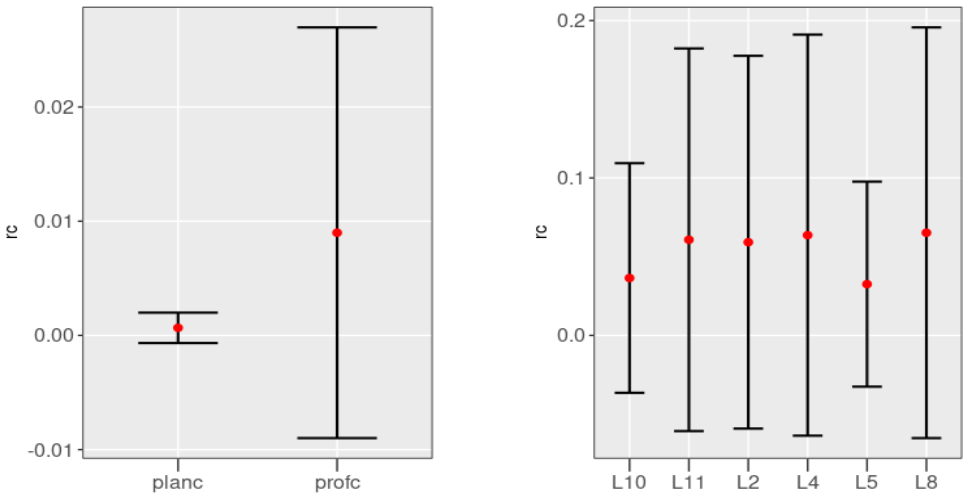

Figure 5.5: Estimated fixed effects of selected covariates on log-surface deformation, with fitted means and corresponding standard error intervals. The left-hand plot displays PlanCμ and ProfCμ and the plot on the right shows the lithology categories.

This complexity likely reflects the diverse geological and structural factors influencing surface deformation near active fault zones, such as lithology, near-surface fault geometry, kinematic variability along strike and interaction between faults at depth (Bloom *et al.*, 2022). As such, this effect is less indicative of a general trend and more suggestive of underlying spatial heterogeneity. Finally, proximity to rivers is associated with greater deformation, which is plausibly due to increased pore water pressure reducing shear strength in slope materials.

Figure 5.7 displays the result of our functional precipitation covariate. The left-hand panel shows a 3D surface plot representing the interaction between the two components of the functional term introduced in Section 5.3.1: the basis matrix for precipitation values, $\mathbf{B}^X$, and the basis matrix for time, $\mathbf{B}^P$, denoted in the figure as `precip.omat` and `precip.tmat`. The resulting surface captures the combined influence of precipitation intensity and timing on surface deformation. We observe both peaks and troughs across the surface, indicating spatial and temporal regions where precipitation exerts a stronger or weaker effect on deformation. The right-hand panel in Figure 5.7 translates this interaction into a spatial contour map, showing the overall effect of precipitation across the study region. The plots show the estimated regression coefficient surface for the functional precipitation predictor, expressed as a smooth function of precipitation magnitude and its temporal basis. The surface represents how precipitation at different magnitudes and times contributes to the modelled log-surface deformation.

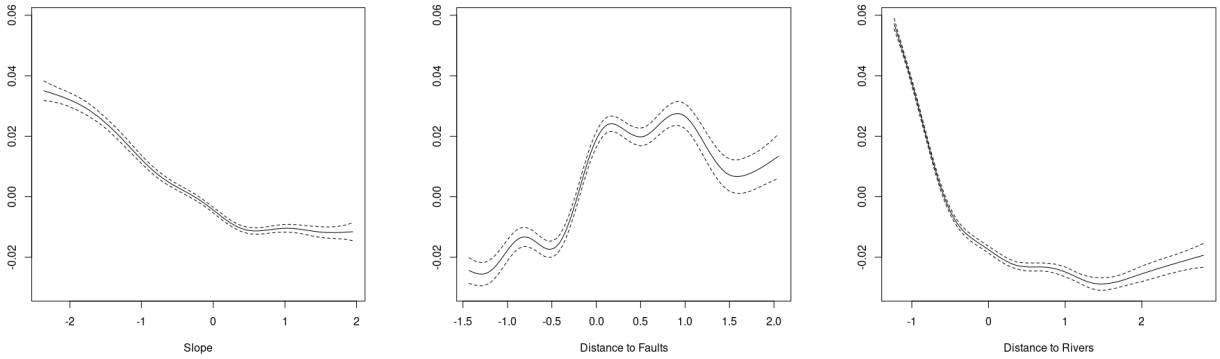

Figure 5.6: Estimated smooth effects of selected covariates on log-surface deformation, with fitted means and corresponding standard error intervals. All covariates were standardised prior to modelling, as reflected in the x-axis scales. The plots highlight non-linear relationships, including increased deformation with steeper slopes and proximity to rivers, and complex variation associated with distance to faults.

The north and north-eastern areas exhibit a stronger precipitation-related deformation signal, while the south-western and eastern areas show a weaker influence. Notably, the high-effect regions (peaks) in the surface plot align with darker areas in the contour map, while low-effect regions (valleys) correspond to lighter contours. This output exemplifies the advantage of using a functional predictor: rather than summarising precipitation with a single coefficient or smooth term, the model yields a nuanced surface that reveals spatially varying and temporally structured effects.

Figure 5.8 displays the spatio-temporal SPDE smooth effect for the first 12 time points across the region in Sichuan, China. Due to the computational cost of fitting the model, the fields are displayed at a relatively coarse resolution for efficiency. Despite this, meaningful spatio-temporal structure is evident. For instance, stronger spatial coherence is observed at time points 2, 6, 8 and 12, whereas time points 1 and 4 exhibit weaker spatial structure. There is no clear temporal trend in the residual spatio-temporal dependence, suggesting that the underlying surface deformation process may not follow a regular spatio-temporal pattern within this domain. This could be influenced by the geometry of the study area, which consists of slope units arranged within a rectangular region, distinct from typical spatial domains such as coastlines or natural landforms. Nonetheless, the spatio-temporal SPDE smoother successfully captures structured residual variation, as demonstrated by the evolving spatial fields over time.

Figure 5.9 presents model diagnostics by comparing observed and fitted values of log-surface defor-

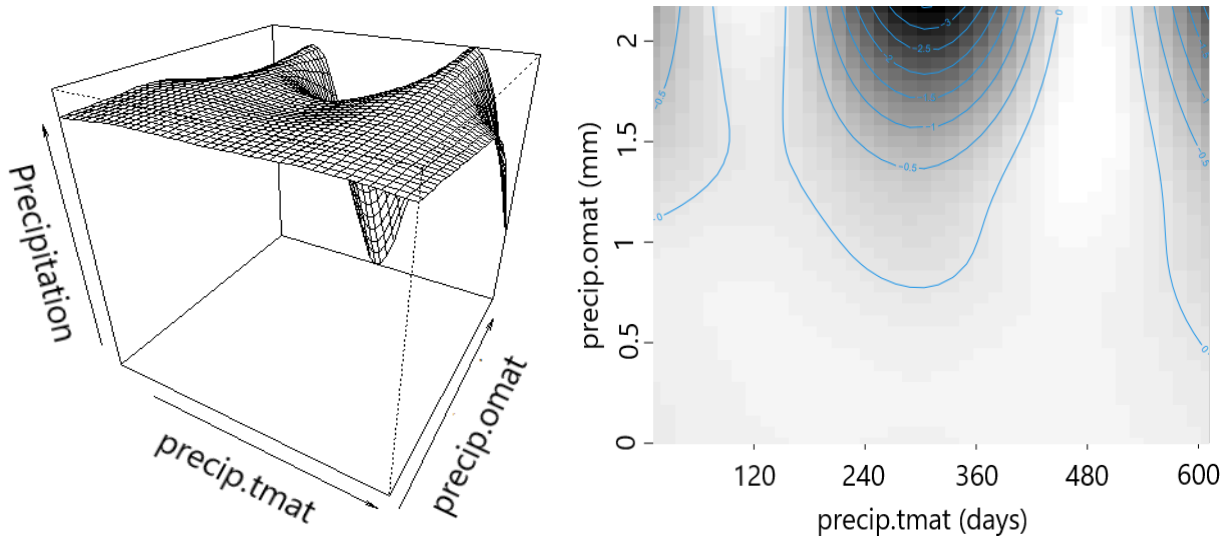

Figure 5.7: Functional effect of precipitation on log-surface deformation. *Left*: 3D surface plot showing the estimated smooth function $F(x, p)$, representing the interaction between precipitation magnitude (x) and time index (p) based on tensor-product basis evaluations $\mathbf{B}^X$ and $\mathbf{B}^P$. *Right*: Spatial contour map displaying the aggregated influence of precipitation across the study region in Sichuan, China. Darker areas indicate regions where precipitation has a stronger effect on deformation, while lighter areas indicate weaker influence.

mation at three selected spatio-temporal instances. These examples are shown for brevity and clarity. Overall, the model captures the general spatial pattern of deformation well. However, discrepancies remain in the extreme values, particularly in regions highlighted in yellow and dark blue at each time point. These areas indicate under- or over-estimation in the tails of the distribution, suggesting that while the model performs well on average, further refinement may be needed to improve accuracy for extreme deformation events.

As a general model diagnostic, Figure 5.10 presents three checks. The left panel shows a QQ plot comparing the empirical quantiles of observed log-surface deformation to those of the fitted values. The model aligns well with the identity line overall, but deviates in the upper tail—consistent with Figure 5.9—indicating that extreme deformation values are underestimated. The central panel displays a histogram of deviance residuals, which shows approximate symmetry around zero, suggesting no systematic bias in the model's predictions. However, the presence of heavy tails hints at a few poorly fitted observations, likely corresponding to extreme values. The right panel is a QQ plot of the residuals against the theoretical quantiles under the fitted log-Gaussian distribution. Mild tail deviations further suggest that the model may not fully capture the distributional structure of the

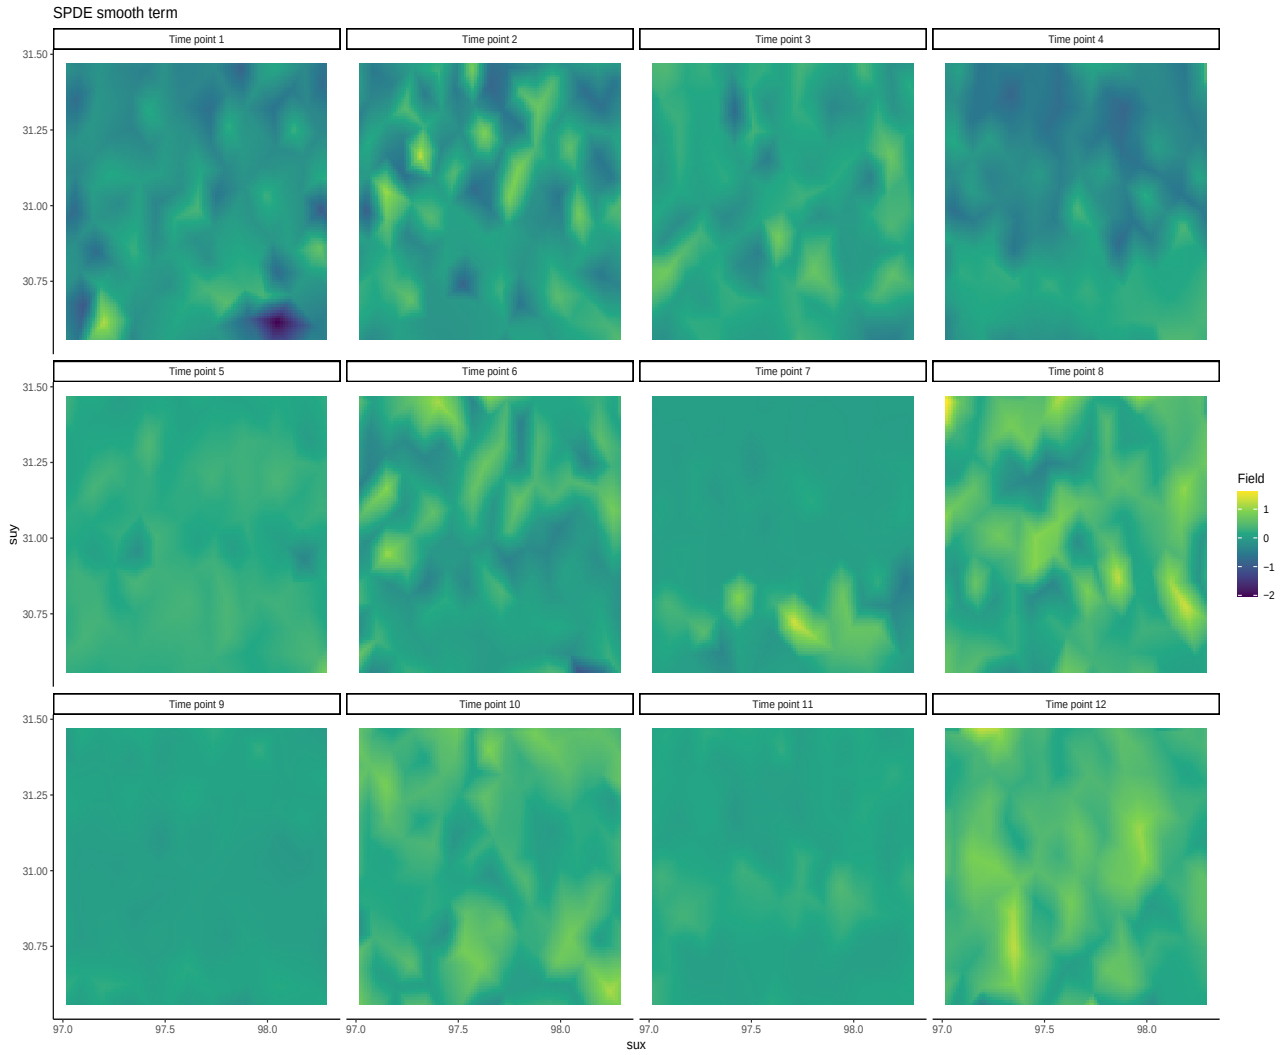

Figure 5.8: Spatio-temporal random fields from the SPDE smoother for the first 12 time points. Each panel shows the spatial variation in the latent field at a given time, capturing residual spatio-temporal dependence not explained by the covariates. Coarser resolution was used for computational efficiency. The spatial patterns vary across time, with stronger structure visible at certain time points (e.g., 2, 6, and 12), indicating dynamic unexplained effects across the region.

residuals, especially at the extremes.

5.6 Conclusion

This work introduces the first implementation of a space-time SPDE smoother within the `mgcv` framework, accompanied by a rigorous mathematical development and practical guidance for implementa-

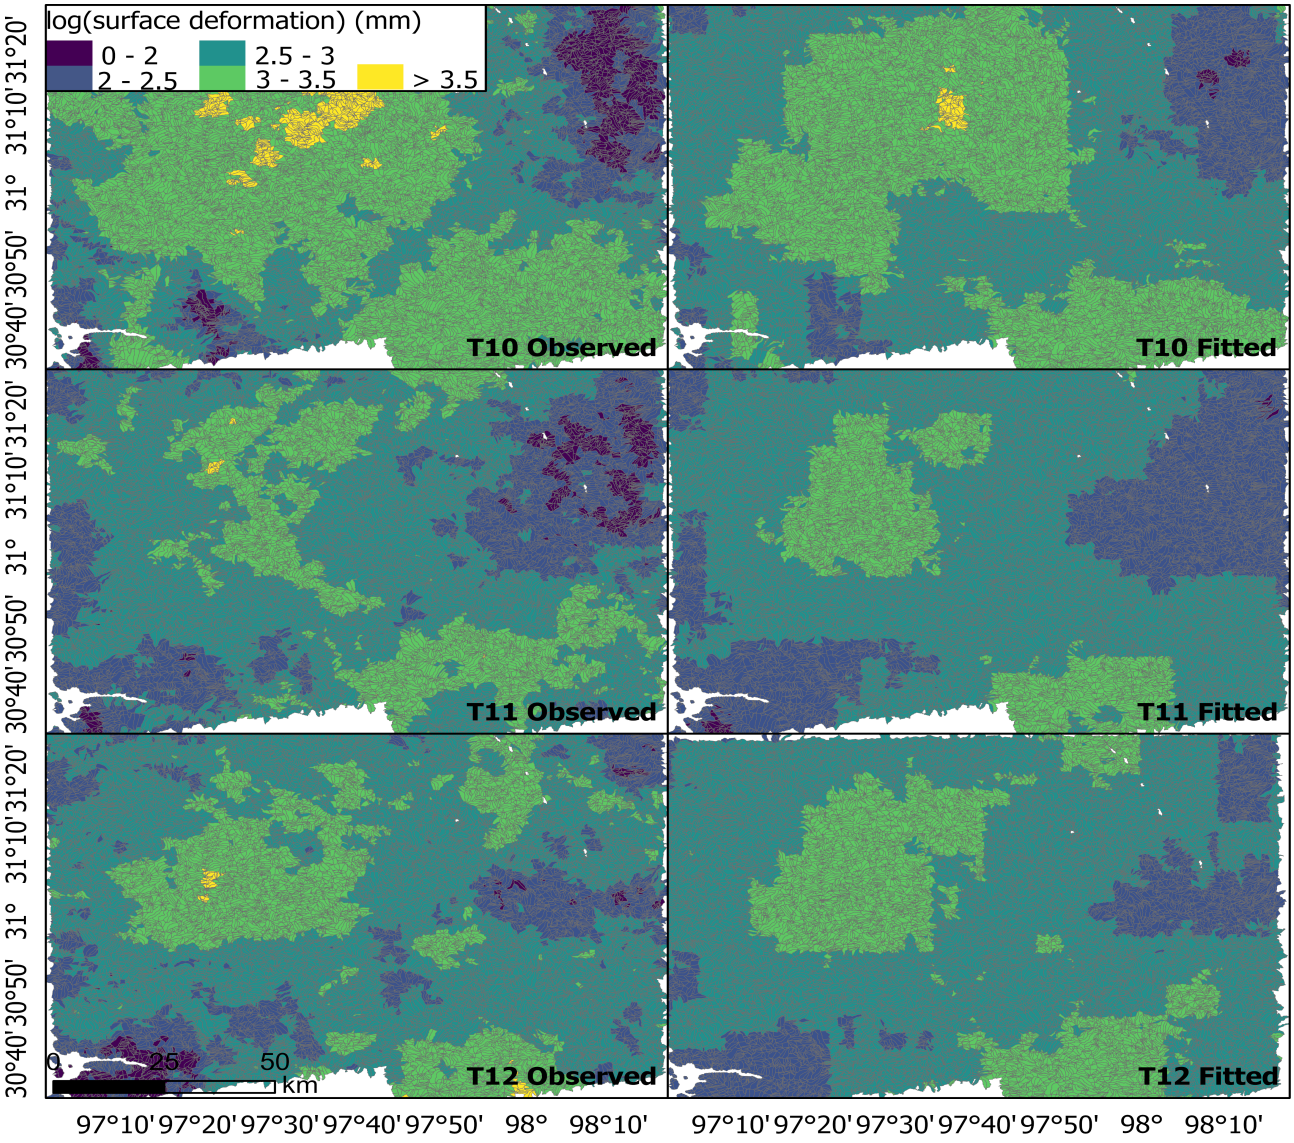

Figure 5.9: Observed (left column) versus fitted values (right column) of log-surface deformation across three points in time for the region.

tion in R. A key innovation lies in enabling the inclusion of a functional covariate—precipitation—within a scalar-on-function regression (SoFR) framework, leveraging the strengths of both functional data analysis and the SPDE approach to spatio-temporal modelling.

Beyond the functional precipitation effect, the model integrates a suite of geographic and geological predictors, combining linear and non-linear components to flexibly model surface deformation. This unified framework yields interpretable results while capturing latent spatio-temporal structure through the SPDE-derived smoother.

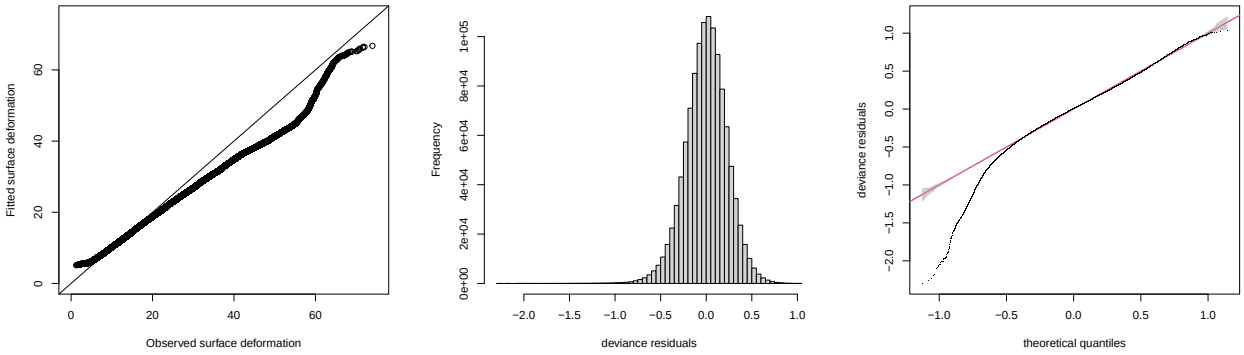

Figure 5.10: Model diagnostics for the log-Gaussian regression model of surface deformation. *Left*: QQ plot comparing observed and fitted values, in original scale, indicating good overall fit but underperformance in the upper tail. *Middle*: Histogram of deviance residuals, approximately symmetric around zero, though with heavy tails suggesting some outlier influence. *Right*: QQ plot of residuals versus theoretical quantiles under the fitted log-Gaussian distribution, showing deviations in the tails, consistent with mild misfit in the extremes.

While the model successfully recovers broad patterns in the deformation field, it shows limitations in capturing extreme values—an issue revealed both in the residual diagnostics and fitted-versus-observed comparisons. This limitation motivates future extensions toward heavier-tailed response models, such as those based on the Generalised Pareto Distribution (GPD), which may better accommodate the tail behaviour characteristic of surface deformation data.

Another avenue for development involves incorporating a non-separable space-time Matérn covariance structure. The current separable formulation simplifies implementation and interpretation but may not fully capture the complex interactions between space and time inherent to landslide processes. Non-separability could provide a more faithful representation of event dynamics and improve model fidelity in hazard contexts.

Chapter 6

Discussion, conclusion, and future work

6.1 Discussion

This thesis provides a comprehensive exploration into landslide hazard assessment modelling, capturing the *where*, *when*, and *how large*. In this chapter, we reflect critically on the choices and modelling approaches made, the limitations of methodologies and data, and potential alternatives or extensions.

6.1.1 On the modelling framework

To begin this discussion, we first evaluate the choice to use a predominantly Bayesian framework. This framework gave us access to flexible hierarchical models, allowing for the inclusion of both covariate effects and latent spatial/spatio-temporal effects in complex models. In addition, the inclusion of prior information is useful for high-risk regions where learning from sparse events is vital. Furthermore, Bayesian models provide full posterior distributions for all model parameters, facilitating the derivation of credible intervals for landslide occurrence probabilities or event sizes, offering richer information for decision-making processes. This is particularly relevant since uncertainty quantification is a critical component in the quantification and communication of landslide risk. Perhaps, the most significant advantage of the Bayesian framework, as utilised throughout this thesis, is that integrated nested Laplace approximation (INLA; Rue *et al.*, 2009) may be used in this framework, which enables computationally efficient and accurate inference for complex latent Gaussian models (LGMs).

INLA has become an increasingly prominent tool in particular in the field of spatio-temporal

modelling (e.g., Rue *et al.* 2009; Lindgren *et al.* 2011; Bakka *et al.* 2018), and continues to evolve to accommodate a broad range of model complexities. One of the primary advantages of INLA over traditional Bayesian inference methods such as Markov Chain Monte Carlo (MCMC) lies in its computational efficiency. While MCMC relies on iterative sampling schemes that can be computationally intensive and time-consuming, particularly in high-dimensional or hierarchical models, INLA employs deterministic approximations based on Laplace approximations and numerical integration to provide accurate posterior estimates at a fraction of the computational cost. For these reasons, we chose INLA for inference on large spatial/spatio-temporal landslide inventories.

Despite its many advantages, the INLA framework has a number of limitations. Its applicability is restricted to models within the class of latent Gaussian models, and unlike more general MCMC-based methods, it cannot be used to estimate parameters for arbitrary models. In particular, INLA may struggle with models involving strong non-Gaussian latent features or complex dependencies among hyperparameters. Moreover, Bayesian inference more broadly is sensitive to the choice of prior distributions. Selecting appropriate priors can be challenging, especially when prior information is scarce or subjective, and conducting thorough sensitivity analyses is often necessary to ensure the robustness of the results.

Furthermore, the interpretability of Bayesian outputs can pose a barrier for practitioners more familiar with frequentist paradigms. While posterior distributions and credible intervals may be more intuitive to some scientists, their understanding of probability is often grounded in point estimates and p-values. This mismatch can complicate the communication of Bayesian results to non-specialist audiences (BiostatMatt, 2025). Lastly, while posterior predictive checks offer a powerful framework for model validation, they can be more complex to implement and difficult to consistently interpret compared to traditional residual diagnostics used in frequentist approaches.

As a final point of evaluation, the direction of this research progressed toward incorporating spatial point process modelling and functional data into landslide hazard models (Chapter 5). This involved extending the framework of susceptibility to accommodate spatial point pattern data, enabling a more detailed representation of event locations and their relationship with environmental covariates. While the inclusion of functional predictors was initially explored within the INLA framework, it proved more practical to employ the `mgcv` package and work to implement a key aspect of the previous framework, the SPDE approach. Consequently, although `R-INLA` was the primary fitting tool for most of the models presented, the shift to an alternative framework underscores the dynamic nature of landslide hazard modelling. It also highlights the importance of methodological flexibility, as the

ability to draw on both approaches - Bayesian and frequentist - proved advantageous.

We next evaluate a key modelling choice: the use of the stochastic partial differential equation (SPDE) approach to capture underlying spatial or spatio-temporal variability in the data. This method provides an efficient approximation to a Gaussian random field with a Matérn covariance structure and is conveniently implemented within the `R-INLA` framework, making it a natural choice for much of the work presented in this thesis.

An alternative considered during the early stages of analysis for the Japanese landslide inventory was the Besag model (Besag, 1974), commonly used for areal data. This model defines spatial dependence through an intrinsic conditional autoregressive (ICAR) structure, in which information is shared between neighbouring regions via a sparse precision matrix derived from a neighbourhood adjacency structure. While computationally efficient and well-suited to modelling local spatial dependencies, the Besag model is limited in its ability to capture long-range spatial correlations. Furthermore, as it is defined over discrete spatial units, it is incompatible with point process models that operate in continuous space, which became a focus of this research. Other possible alternatives include spatial basis function representations such as fixed-rank kriging (Cressie and Johannesson, 2008; Gabriel *et al.*, 2016) or Gaussian predictive processes (Banerjee *et al.*, 2008); however, these approaches were not implemented in this work due to the demonstrated adequacy of the SPDE approach and its established validity in point process models.

Lindgren *et al.* (2011) established an explicit connection between Matérn Gaussian random fields (GRFs) and a class of SPDEs, enabling computationally efficient representations through Gaussian Markov random fields (GMRFs) with sparse precision matrices. The parameters of the Matérn covariance function govern key characteristics of the field, including its smoothness, spatial range, and marginal variance. While alternative covariance structures - such as rational quadratic or non-stationary kernels - can be specified within a full Gaussian process framework using MCMC, they generally do not yield sparse precision matrices, making them less computationally efficient and more complex to implement.

In terms of computational resources, the Hurdle model in Chapter 2 had a run-time in the region of less than ten minutes with the associated unified exceedance probabilities also running quickly (less than ten minutes) despite the posterior sampling undertaken. The susceptibility model and the approximation to the LGCP fitted in Chapter 3 also had a run-time in the region of less than ten minutes. The simulation studies of Chapter 4 ranged in run-times to less than one minute (coarse mesh and smooth covariate) to within the hour (fine mesh and fine-scale covariate), while the multiple

marked LGCPs fitted to local subsets of Japan all ran within ten minutes. The space-time model of Chapter 5 had to be run through an external Geospatial Computing platform (Crib) and the model took between 24 and 48 hours to run. This is because `mgcv` deals with dense matrices rather than sparse, and as such, the space-time SPDE smoother involved multiple Kronecker products of dense matrices.

As outlined previously, the research evolved toward spatial point process modelling, primarily due to the availability of precise landslide occurrence coordinates. This was the case for both the debris flow data across Scotland provided by the British Geological Survey (BGS) and the Japanese landslide inventory. In both datasets, covariates were defined over irregular spatial units—slope units (SUs) for Scotland and catchments for Japan - allowing point occurrences to be aggregated into counts per unit where appropriate. This aggregation facilitated the development of a log-Gaussian Cox process (LGCP) approximation for the BGS dataset. While Poisson count models offer advantages in terms of computational efficiency and ease of interpretation, spatial point process models operate over continuous space, providing higher spatial resolution and more detailed inference. Additionally, point process models support extensions to include 'marks', such as landslide size, enabling joint modelling of occurrence and size. Given the aim of characterising landslide hazard in terms of both frequency and intensity, the point process framework was deemed more suitable for this research in comparison to joint models at the spatial units.

We would like to highlight the comparisons; between the point process approach to other spatial modelling used for landslides in the literature and throughout this research, and the representation and approximation of the Gaussian random field through the SPDE (and INLA) to other options - e.g. using a covariance function to characterise the fields and MCMC to estimate it (Simpson *et al.*, 2012).

6.1.2 The Hurdle model

We now evaluate the modelling choices made for the Hurdle model proposed for landslide occurrence and size on the Island of Dominica after Hurricane Maria in 2017, outlined in Chapter 2. Notably, here we model two aspects of landslide hazard: the spatial distribution and 'intensity', characterised by planimetric extent. The lack of a temporal aspect was due to the limited access to long-time series of landslide inventories triggered in response to hurricanes across the Island. However, for the purpose of this work, short-time series of the 2017 landslides would have been sufficient for a unified landslide hazard assessment. A second limitation of these data is the lack of a rainfall predictor. This was

due to the weather stations failing during Maria and only a few providing reliable data (Dominica Meteorological Service, 2019). From the stations that were able to record rainfall measurements, over 500mm were recorded, with amounts increasing from 5 : 00pm on the 18^{th} (Dominica Meteorological Service, 2019), this endangered the upper catchments to be susceptible to debris flows that could be triggered with rainfall thresholds substantially lower than before the hurricane (ReliefWeb, 2018). Lastly, it was noted that out of the landslides that occurred on Dominica, 8,576 were debris slides, 1,010 were debris flows and 374 were rock falls (ReliefWeb, 2018). If this classification had been available, it would have been possible to model the three landslide types separately, offering greater insight into which covariates influence each type and where on the Island each class is most likely to occur following a hurricane. Thus making it clearer as to which safety measures should be prioritised and where. Another consideration is whether the log-Gaussian distribution was the most suitable choice for modelling landslide size. It was selected based on the approximately bell-shaped nature of the observed data; however, given the interest in exceedance probabilities and the potential presence of heavy tails, further exploration of alternative distributions - such as the generalised Pareto - may be warranted to better capture the underlying variability and extremes.

The model performance measures undertaken in this study were a receiving operator characteristic (ROC) curve for the predicted probability of landslide occurrence, a probability integral transform (PIT) histogram and an observed versus predicted values scatter-plot for the log-landslide sizes. These measures were chosen for their robustness to class imbalance (ROC curve-many landslide absences in comparison to presences), checking for over- or under-prediction of the landslide log-sizes, and to visualise the comparison between observed and fitted values (PIT Histogram and observed versus fitted plot).

6.1.3 The BGS data

We now evaluate the modelling choices made for the updated landslide susceptibility model for the debris flow data across Scotland by the BGS and the LGCP extension, outlined in Chapter 3. The development of the statistical methodology for this chapter was, to some extent, constrained by the nature of the available data. However, the application presented in this chapter is noteworthy as it is the first data-driven landslide susceptibility model proposed for a world-leading independent research organisation. Available were the debris flow point coordinates and a binary presence/absence indicator for the SU, with stakeholder interest focused solely on landslide susceptibility. Consequently, this chapter addresses only one component of landslide hazard - its spatial distribution - while the

exclusion of temporal and intensity dimensions reflects the constraints imposed by the nature of the project. Although not all debris flow points had associated temporal information at the time of data collection, restricting the analysis to only those with complete timestamps would have substantially reduced the inventory size, resulting in a model that would not realistically represent Scottish debris flows.

A secondary data limitation is that mean annual averages of the maximum daily precipitation over a 20-year period were used as a rainfall predictor. This single-statistic summary is a reduction of a much richer and temporally detailed dataset. Ideally, precipitation as a continuous function would offer critical insight into ground conditions, saturation properties, and rainfall intensity patterns - factors that are central to understanding and modelling the conditions leading to landslide initiation. While this limitation constrained the scope of the current chapter, it motivated a valuable line of inquiry for subsequent research presented in Chapter 5: *how can precipitation be incorporated as a functional predictor within the available modelling frameworks?*

The decision to present an LGCP extension to the landslide susceptibility arose from the availability of the debris flow point coordinates, and a personal interest to work with a more complex model. This choice certainly created a domino effect. A significant amount of time was spent working with the BGS data and grappling with the *mess with the mesh* issue (see Chapter 4), before realising that this was not a problem isolated to just this dataset. Rather, it was a common challenge often overlooked when dealing with high-resolution data and the SPDE approach. Due to time constraints and limited success in fully implementing the LGCP, we ultimately chose to fit an LGCP approximation using Poisson regression instead. While this approach allowed us to provide more detailed information on debris flow patterns across Scotland, it was not the complex model initially desired. Nevertheless, this path led to a deeper investigation into the *mess with the mesh* issue, which ultimately paved the way for successful point process model fitting with high-resolution landslide point pattern data.

The model performance measures undertaken for this study were a ROC curve and ROC curves for a 10-fold cross-validation (CV) procedure, applied to the predicted probabilities of debris flow occurrence. Additionally, we employed an observed versus fitted debris flow count scatter plot, as well as observed versus fitted plots from the 10-fold CV implementation for evaluating the predicted intensity of the debris flow point pattern (via the LGCP approximation). The strengths and limitations of ROC curves and observed versus fitted plots have been discussed earlier; however, it is worth noting here that, in this context, the use of a goodness-of-fit measure with broad interpretability proved advantageous.

In the cross-validation (CV) implementations, the dataset was randomly partitioned into ten subsets. For each iteration, the model was trained on nine of these subsets and subsequently used to predict over the one subset that was left out. This process was repeated such that each subset served as the validation set exactly once. Ten-fold CV provides a robust assessment of model generalisability, as the model is evaluated on unseen data in each iteration. However, a key limitation of ten-fold CV is its underlying assumption that data points are independent and identically distributed (i.i.d.). In the context of landslide hazard modelling, this assumption is often violated due to spatial autocorrelation. This spatial dependence can lead to overly optimistic estimates of model performance. Additionally, the debris flow dataset used in this study is characterised by a pronounced presence/absence imbalance, with a large number of absences and relatively few presences. When folds are selected randomly, this imbalance can introduce bias into the model evaluation. This limitation also motivated the use of an LGCP, since in spatial point process modelling, the absence of events across the space carries as much weight as the locations where events do occur. Another drawback of ten-fold CV is that it does not assess the model's ability to generalise to entirely new geographic regions - a capability that is particularly relevant for national-scale applications. As an improvement, a spatial cross-validation technique such as leave-group-out cross-validation (LGOCV) explicitly accounts for spatial structure in the data and provides a more realistic assessment of predictive performances across space (Adin *et al.*, 2024).

Validation measures for spatial point processes are complex due to their continuity in space: when a randomly selected area of the study region is used for prediction, it may either contain all, some, or none of the observed points. Removing a point, or several, alters the pattern and consequently the data. With geo-statistical data, points can be removed and the model fitted without these, but if points are removed in a point pattern, the nature of the observation is changed (one pattern is one observation). Therefore, an effective assessment for spatial point processes is a Bayesian approach to residual analysis (Baddeley *et al.*, 2005). This proposes approaches to evaluating Gibbs processes however, rather than LGCPs. During the initial stages of the BGS analysis, residual analysis was employed to compare model fits under varying combinations of covariates and the inclusion of a spatial random effect. However, with the emergence of the *mess with the mesh* issue, the modelling approach was simplified to a Poisson regression. This change enabled the use of a ten-fold cross-validation (CV), which, for the sake of consistency with the susceptibility model, was adopted as the primary validation method.

6.1.4 The *mess with the mesh*

We now evaluate the modelling choices made in relation to the *mess with the mesh* analysis presented in Chapter 4. This chapter highlights a common issue that arises when implementing the SPDE approach to point process models using high-resolution spatial data, illustrated through an application in landslide hazard modelling. Specifically, we present a spatial marked LGCP model to describe the spatial distribution and associated intensity of landslides across Japan. The primary objective of this chapter was to draw attention to the mesh-related issue and to demonstrate how it can emerge under certain data conditions; thus, the application itself served a largely illustrative role rather than being the central focus. As such, the omission of a temporal dimension - while it could have further enriched the analysis - was not considered critical, given that relevant temporal data were not available. Instead, the analysis was used to effectively demonstrate the nature of the *mess with the mesh* problem. Similarly, the precipitation predictor was again simplified to a maximum rainfall average statistic. While this is a limitation from an application perspective, it is consistent with the illustrative purpose of the study. In this vein, formal goodness-of-fit measures were not comprehensively undertaken. Instead, model adequacy was evaluated using QQ plots comparing the distribution of observed versus fitted landslide sizes across subsets, as well as a plot of mean landslide intensity overlaid with observed landslide point locations, accompanied by a summary of observed versus estimated landslide counts by subset. These diagnostic tools were deemed sufficient to demonstrate the feasibility and effectiveness of implementing a spatial point process model with high-resolution covariate data using the SPDE approach.

An important point of discussion concerns the design of the simulation studies used to illustrate the mesh-related issue. Simulation study one demonstrates the sensitivity of model results to the definition of integration points, showing how slight shifts in the domain of integration can lead to substantial variation in the estimated regression coefficients for model covariates. This was evaluated using three point process models, each with an intensity function explained by a single covariate, differing in spatial range from fine (short-range) to coarse (long-range). A single covariate was intentionally selected to isolate the effects of shifting integration domains, avoiding the additional complexity introduced by a spatial random effect. However, the inclusion of covariates with differing spatial ranges may have inadvertently introduced complexity beyond the intended focus of domain shifts.

Simulation study two demonstrates how the inclusion of a spatial random effect can further amplify the variability in regression coefficient estimates. However, as shown in Simulation study one,

the underlying issue does not stem from the inclusion of the spatial random effect itself. Rather, it arises from the definition of the integration points, which are fundamental to the likelihood formulation. All model components - both covariate effects and the spatial random field - are evaluated at these integration points. This study underscores the extent to which parameter estimates for the spatial field can vary depending on the resolution of the mesh used to approximate the domain. In particular, it highlights the sensitivity of spatial model inference to mesh construction, reinforcing the importance of aligning mesh resolution with both the scale of spatial variation and the underlying data characteristics.

A final point to note is that the marked LGCP model presented is not a global model. Due to the nature of the *mess with the mesh* issue (we are working with high-resolution covariate data) and in the interest of computational efficiency, the modelling approach involved fitting local models to subsets of the data, from which a broader, global effect was inferred. While this approach is not ideal for representing landslide hazard at a national scale, it was a reasonable compromise given that, as noted above, the primary objective was not the application itself but rather to highlight the challenges of using high-resolution spatial covariates within the SPDE framework.

6.1.5 On the space-time SPDE smoother

Here, we critically evaluate the modelling decisions made in Chapter 5, where a spatio-temporal model was developed to analyse surface deformation across a region in Sichuan, China. This work develops a space-time SPDE smoothing term within the `mgcv` framework and employs rainfall as a functional predictor through scalar-on-function regression (SoFR). The motivation for constructing a custom SPDE smoothing term in `mgcv` stemmed from the desire to integrate rainfall as a functional predictor within a frequentist modelling framework, while also retaining elements of Bayesian spatial modelling. Notably, the choice to adopt functional data analysis and SoFR - rather than adapt the INLA framework to accommodate the functional nature of the covariate - was informed by initial attempts to incorporate functional rainfall within INLA. These attempts consistently yielded only a single, time-invariant estimate of the effect, rather than a continuous functional representation across time and space. As a result, we redirected our efforts toward the `mgcv` framework, where we implemented the space-time SPDE smoother.

We also explored alternative spatio-temporal modelling approaches using INLA, developing models with both separable and non-separable space-time covariance structures. In these models, rainfall was included as a covariate via a random effect, modelled using a second-order random walk (RW2; see

Section 1.3). These exploratory models provided useful insight into the behaviour of space-time fields under different modelling assumptions and highlighted the contrast between incorporating rainfall as a smooth functional predictor versus a random non-linear effect. Although these models are not discussed in detail in the final version of the thesis, they contributed to the broader evaluation of modelling strategies for surface deformation and informed the decision to pursue the `mgcv`-based approach.

Due to the structure of the custom `smooth.construct` function developed for the space-time SPDE smooth term - specifically, its reliance on multiple penalty matrices constructed via Kronecker products - the model is computationally intensive. As a result, although surface deformation observations and rainfall data were available for the period 2017 to 2021, it was more practical to model a shorter temporal window, focusing on only nineteen months of data. To further accommodate the computational demands and dataset size, the `bam` function (designed for handling large datasets) was used in place of `gam` within the `mgcv` framework.

6.2 Conclusion

This thesis presents a range of statistical methodologies for landslide hazard assessment, with a primary focus on Bayesian hierarchical modelling. It addresses both methodological and practical challenges associated with capturing latent spatial structures characteristic of landslide-prone environments. These structures are represented through carefully selected, domain-relevant covariates, and, in the final case study, extended to incorporate a functional process that varies across both space and time.

This thesis provides a comprehensive overview of landslide hazard modelling, outlining the progression of statistical models. These models include susceptibility (presence/absence), the joint hurdle model, point process approaches, spatio-temporal models, and accounting for landslide distributions through various likelihood functions. Alongside this exploration is an account for modelling challenges that may arise with high-resolution covariate data, such as the landslide inventories available, and a new avenue into converting the SPDE approach for use in a frequentist framework.

Throughout the work, we build on the definition of 'landslide hazard' by Guzzetti *et al.* (1999); at first considering the spatial distribution and associated intensity, and improving on the modelling aspects such that a full assessment of landslide hazard may be presented in the final case study-that of the *where*, *when* and *how large* of a landslide process (ground surface deformation). The case

studies included in this work each bring a unique aspect to the assessment of landslide hazard. The Hurdle model enables the simultaneous estimation of both the probability of landslide occurrence and its expected planimetric extent, supporting probabilistic hazard maps with quantifiable uncertainty and exceedance probabilities. Applied to post-hurricane data from the Island of Dominica, this study demonstrates the potential of landslide hazard assessment in a post-disaster risk scenario.

The second applied contribution involves the development of a data-driven debris flow susceptibility model for Scotland, in collaboration with the British Geological Survey (BGS). A comparison with the previous heuristic GeoSure model highlights the practical benefits of probabilistic methods in operational hazard mapping, particularly in assigning risk probabilities to road and traffic networks. The log-Gaussian Cox process (LGCP) extension revealed important implementation challenges when applying point process models with the SPDE approach to high-resolution data. These challenges - centred on the specification of the domain of integration - are systematically explored through simulation studies and resolved through recommended guidance applicable to broader spatial modelling contexts. The relevance of this work is highlighted through its application to a subset of Japan's landslide inventory, extending the framework to a marked spatial point process to model this mark, the landslide size.

The final contribution of this thesis marks a shift towards a frequentist functional data approach, designed to more effectively capture the complexity of precipitation as a key predictor. To support this, a novel space-time SPDE smoother was developed within the `mgcv` framework, integrating a core element of the INLA methodology - the SPDE approach for modelling spatial and temporal variation - into a frequentist setting. The resulting case study, which models surface deformation across a region in Sichuan, China over a nineteen-month period, effectively concludes the thesis's exploration on the *where*, *when*, and *how large* of landslides.

In summary, this thesis has integrated methodological development with landslide application, presenting a novel statistical tool and providing broader insights into the modelling and assessment of landslide hazard through a range of landslide inventories.

6.3 Future work

In this section, we outline possible extensions to the current work.

Perhaps the most crucial potential extension for this work is in the validation of the point process models, which, as previously noted, is not straightforward. Cronie *et al.* (2024) propose a statistical

framework for validating point processes by thinning the process into training and validation sets. However, further investigation is needed to assess whether this approach preserves the notion of 'sample size,' since technically a point process is a single realisation. As the thinning process involves splitting the observed pattern into separate training and validation sets, this may substantially alter the properties of the point process and thus affect the validity of model assessment.

Another avenue for potential improvement is exploring different link functions for the data. Throughout this thesis, we modelled the mean of the response variable, μ, as a function of a linear predictor, η, typically using canonical link functions such as the logit link for Bernoulli models (e.g., in the Hurdle model and BGS susceptibility) and the identity link for continuous outcomes such as landslide log-sizes and surface deformation. However, in cases of imbalanced classification—such as rare landslide occurrence - the use of alternative link functions like the complementary log-log (cloglog) link may be more appropriate, as it accommodates asymmetric response behaviour and is better suited for rare event modelling. More generally, custom link functions can be introduced to reflect skewness or other characteristics of the data, provided they remain invertible and differentiable (McCullagh, 2019). This flexibility is particularly relevant in latent variable models, where skewness or heavy tails may be more naturally captured in the latent process.

A secondary direction for future research involves the implementation of spatial cross-validation techniques, particularly leave-group-out cross-validation (LGOCV) for the surface deformation FGAM, as well as Bayesian residual analysis for the marked LGCP applied to landslide occurrence and size. These approaches would allow for a more rigorous evaluation of model robustness beyond hazard estimation, providing diagnostic measures that are both methodologically sound and practically interpretable.

More specifically, an important improvement to the BGS landslide susceptibility model would be to recover the subset of debris flow data that includes temporal information and use it to develop a spatio-temporal marked point process model. This would allow for a full log-Gaussian Cox process (LGCP) implementation rather than relying on a Poisson approximation, while incorporating the insights gained from the *mess with the mesh* study to ensure appropriate integration over space and time. For the spatio-temporal FGAM for surface deformation, defining a non-separable space-time SPDE smoother would be the next natural extension of this work. Separability in the space-time covariance structure is more convenient for implementation but for application purposes, particularly landslide modelling, observations in space and time are likely to be dependent and non-separable in their influence on occurrence. Thus, deriving a differential operator corresponding to the non-

separable covariance function represents a promising direction for enhancing this work.

Non-separable space-time models in `R-INLA` have been introduced by Bakka *et al.* (2020) through a diffusion-based extension of the Matérn field, known as the Diffusion-based Extended Matérn Field (DEMF). This approach constructs a stochastic process grounded in diffusion equations, such as the heat equation, leading to inherent space-time covariance non-separability. The DEMF model is defined by the equation

$$\left(\gamma_t \frac{\delta}{\delta t} + L_s^{\alpha_s/2}\right)^{\alpha_t} u(\mathbf{s}, t) = \delta\varepsilon_{\mathrm{Q}}(\mathbf{s}, t),$$

where $u(\mathbf{s}, t)$ is the spatio-temporal process (typically appearing in the linear predictor), $\delta\varepsilon_{\mathrm{Q}}(\mathbf{s}, t)$ denotes temporally uncorrelated but spatially structured noise, and $L_s = \gamma_s^2 - \Delta$ is the spatial differential operator. The parameters (α_s, α_t) control the smoothness in space and time, respectively, while (γ_s, γ_t) are positive scale parameters. The corresponding differential operator is

$$D_{nst} = \left(\gamma_t \frac{\delta}{\delta t} + L_s^{\alpha_s/2}\right)^{\alpha_t}.$$

From here, one could calculate $\int_T \int_S \{D_{nst} f\}^2 d\mathbf{s} dt$ as the smoothing penalty for implementation in `mgcv`.

These extensions provide a theoretical insight into the further capabilities of landslide hazard modelling, and succinctly conclude the extent of this research.

Bibliography

ACAPS (2018) Dominica; the impact of hurricane Maria. https://www.acaps.org/sites/acaps/files/products/files/20180131_acaps_disaster_profile_dominica_v2.pdf.

Adin, A., Krainski, E. T., Lenzi, A., Liu, Z., Martínez-Minaya, J. and Rue, H. (2024) Automatic cross-validation in structured models: is it time to leave out leave-one-out? *Spatial Statistics* **62**, 100843.

Alvioli, M., Marchesini, I., Reichenbach, P., Rossi, M., Ardizzone, F., Fiorucci, F. and Guzzetti, F. (2016) Automatic delineation of geomorphological slope units with r. slopeunits v1. 0 and their optimization for landslide susceptibility modeling. *Geoscientific Model Development* **9**(11), 3975–3991.

Amato, G., Eisank, C., Castro-Camilo, D. and Lombardo, L. (2019) Accounting for covariate distributions in slope-unit-based landslide susceptibility models. a case study in the alpine environment. *Engineering geology* **260**, 105237.

Auguie, B. (2022) *gridExtra: miscellaneous functions for "grid" graphics*. R package version 2.3.

Azarafza, M., Azarafza, M., Akgün, H., Atkinson, P. M. and Derakhshani, R. (2021) Deep learning-based landslide susceptibility mapping. *Scientific Reports* **11**(1), 24112.

Bachl, F. E., Lindgren, F., Borchers, D. L. and Illian, J. B. (2019) inlabru: an R package for Bayesian spatial modelling from ecological survey data. *Methods in Ecology and Evolution* **10**, 760–766.

Baddeley, A., Bárány, I. and Schneider, R. (2007) Spatial point processes and their applications. *Stochastic Geometry: Lectures Given at the CIME Summer School Held in Martina Franca, Italy, September 13–18, 2004* pp. 1–75.

Baddeley, A., Turner, R., Møller, J. and Hazelton, M. (2005) Residual analysis for spatial point processes (with discussion). *Journal of the Royal Statistical Society Series B: Statistical Methodology* **67**(5), 617–666.

Bagchi, R. and Illian, J. B. (2015) A method for analysing replicated point patterns in ecology. *Methods in Ecology and Evolution* **6**(4), 482–490.

Bakka, H., Krainski, E., Bolin, D., Rue, H. and Lindgren, F. (2020) The diffusion-based extension of the Matérn field to space-time .

Bakka, H., Rue, H., Fuglstad, G.-A., Riebler, A., Bolin, D., Illian, J., Krainski, E., Simpson, D. and Lindgren, F. (2018) Spatial modeling with R-INLA: A review. *Wiley Interdisciplinary Reviews: Computational Statistics* **10**(6), e1443.

Ballantyne, C. K. (1986) Landslides and slope failures in Scotland: a review. *Scottish Geographical Magazine* **102**(3), 134–150.

Banerjee, S., Gelfand, A. E., Finley, A. O. and Sang, H. (2008) Gaussian predictive process models for large spatial data sets. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **70**(4), 825–848.

Bengtsson, H. (2024) *matrixStats: functions that apply to rows and columns of matrices (and to vectors)*. R package version 1.4.1.

Besag, J. (1974) Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society: Series B (Methodological)* **36**(2), 192–225.

Besag, J., York, J. and Mollié, A. (1991) Bayesian image restoration, with two applications in spatial statistics. *Annals of the Institute of Statistical Mathematics* **43**, 1–20.

Bewick, V., Cheek, L. and Ball, J. (2004) Statistics review 13: receiver operating characteristic curves. *Critical Care* **8**, 1–5.

Bhuyan, K., Meena, S. R., Nava, L., van Westen, C., Floris, M. and Catani, F. (2023) Mapping landslides through a temporal lens: an insight toward multi-temporal landslide mapping using the u-net deep learning model. *GIScience & Remote Sensing* **60**(1), 2182057.

BiostatMatt (2025) Bayesian vs. frequentist intervals: Which are more natural to scientists? https://biostatmatt.com/archives/1812.

Bivand, R., Gómez-Rubio, V. and Rue, H. (2015) Spatial data analysis with R-INLA with some extensions.

Bivand, R., Keitt, T. and Rowlingson, B. (2023) *rgdal: Bindings for the 'Geospatial' Data Abstraction Library*. R package version 1.6-10.

Bivand, R. and Lewin-Koh, N. (2023) *maptools: Tools for Handling Spatial Objects*.

Bivand, R. and Rundel, C. (2023) *rgeos: Interface to Geometry Engine – Open-Source (GEOS)*. R package version 0.6-3.

Blangiardo, M., Cameletti, M., Baio, G. and Rue, H. (2013) Spatial and spatio-temporal models with R-INLA. *Spatial and Spatio-temporal Epidemiology* **4**, 33–49.

Bloom, C. K., Howell, A., Stahl, T., Massey, C. and Singeisen, C. (2022) The influence of off-fault deformation zones on the near-fault distribution of coseismic landslides. *Geology* **50**(3), 272–277.

Bolin, D. and Kirchner, K. (2020) The rational SPDE approach for Gaussian random fields with general smoothness. *Journal of Computational and Graphical Statistics* **29**(2), 274–285.

Van den Bout, B., Lombardo, L., Chiyang, M., van Westen, C. and Jetten, V. (2021) Physically-based catchment-scale prediction of slope failure volume and geometry. *Engineering Geology* **284**, 105942.

Box, G. and Wilson, K. (1951) On experimental attainment of optimum conditions, ronal stat. *Soc., B* **13**.

Bren, L. (2023) The basics of catchment hydrology. In *Forest Hydrology and Catchment Management: An Australian Perspective*, pp. 1–27. Springer.

Brenner, S. C. (2008) *The Mathematical Theory of Finite Element Methods*. Springer.

Budimir, M., Atkinson, P. and Lewis, H. (2015) A systematic review of landslide probability mapping using logistic regression. *Landslides* **12**, 419–436.

Cameletti, M., Lindgren, F., Simpson, D. and Rue, H. (2013) Spatio-temporal modeling of particulate matter concentration through the SPDE approach. *AStA Advances in Statistical Analysis* **97**, 109–131.

Carrara, A., Cardinali, M., Detti, R., Guzzetti, F., Pasqui, V. and Reichenbach, P. (1991) GIS techniques and statistical models in evaluating landslide hazard. *Earth Surface Processes and Landforms* **16**(5), 427–445.

Chang, J.-M., Chen, H., Jou, B. J.-D., Tsou, N.-C. and Lin, G.-W. (2017) Characteristics of rainfall intensity, duration, and kinetic energy for landslide triggering in taiwan. *Engineering Geology* **231**, 81–87.

Chang, Z., Huang, J., Huang, F., Bhuyan, K., Meena, S. R. and Catani, F. (2023) Uncertainty analysis of non-landslide sample selection in landslide susceptibility prediction using slope unit-based machine learning models. *Gondwana Research* **117**, 307–320.

Chigira, M. and Yokoyama, O. (2005) Weathering profile of non-welded ignimbrite and the water infiltration behavior within it in relation to the generation of shallow landslides. *Engineering Geology* **78**(3-4), 187–207.

Choiruddin, A., Coeurjolly, J.-F. and Waagepetersen, R. (2021) Information criteria for inhomogeneous spatial point processes. *Australian & New Zealand Journal of Statistics* **63**(1), 119–143.

Cigna, F., Bateson, L. B., Jordan, C. J. and Dashwood, C. (2014) Simulating sar geometric distortions and predicting persistent scatterer densities for ers-1/2 and envisat c-band sar and insar applications: Nationwide feasibility assessment to monitor the landmass of great britain with sar imagery. *Remote Sensing of Environment* **152**, 441–466.

Corominas, J., van Westen, C., Frattini, P., Cascini, L., Malet, J.-P., Fotopoulou, S., Catani, F., Van Den Eeckhaut, M., Mavrouli, O., Agliardi, F. *et al.* (2014) Recommendations for the quantitative analysis of landslide risk. *Bulletin of Engineering Geology and the Environment* **73**, 209–263.

Crainiceanu, C. M., Goldsmith, J., Leroux, A. and Cui, E. (2024) *Functional Data Analysis with R*. Volume 174 of *Monographs on Statistics and Applied Probability*. Boca Raton: Chapman and Hall/CRC. ISBN 9781032244716.

Cressie, N. and Johannesson, G. (2008) Fixed rank kriging for very large spatial data sets. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **70**(1), 209–226.

Cronie, O., Moradi, M. and Biscio, C. A. (2024) A cross-validation-based statistical theory for point processes. *Biometrika* **111**(2), 625–641.

Crozier, M. (2005) Multiple-occurrence regional landslide events in New Zealand: hazard management issues. *Landslides* **2**(4), 247–256.

Dambly, L. I., Isaac, N. J., Jones, K. E., Boughey, K. L. and O'Hara, R. B. (2023) Integrated species distribution models fitted in INLA are sensitive to mesh parameterisation. *Ecography* **2023**(7), e06391.

Di Napoli, M., Carotenuto, F., Cevasco, A., Confuorto, P., Di Martire, D., Firpo, M., Pepe, G., Raso, E. and Calcaterra, D. (2020) Machine learning ensemble modelling as a tool to improve landslide susceptibility mapping reliability. *Landslides* **17**(8), 1897–1914.

Diggle, P. J., Moraga, P., Rowlingson, B. and Taylor, B. M. (2013) Spatial and spatio-temporal log-Gaussian Cox processes: extending the geostatistical paradigm .

Dominica Meteorological Service (2019) Hurricane maria: Hydro-meteorological impact on Dominica. Technical report, Dominica Meteorological Service. Accessed: 12 June 2025.

Dong, A., Dou, J., Fu, Y., Zhang, R. and Xing, K. (2023) Unraveling the evolution of landslide susceptibility: a systematic review of 30-years of strategic themes and trends. *Geocarto International* **38**(1), 2256308.

Douglas Nychka, Reinhard Furrer, John Paige and Stephan Sain (2021) fields: Tools for spatial data. R package version 16.3.1.

Dovers, E., Popovic, G. C. and Warton, D. I. (2024) A fast method for fitting integrated species distribution models. *Methods in Ecology and Evolution* **15**(1), 191–203.

Fang, Z., Wang, Y., van Westen, C. and Lombardo, L. (2024) Landslide hazard spatio-temporal prediction based on data-driven models: estimating where, when and how large landslide may be. *International Journal of Applied Earth Observation and Geoinformation* **126**, 103631.

Fayne, J. V., Ahamed, A., Roberts-Pierel, J., Rumsey, A. C. and Kirschbaum, D. (2019) Automated satellite-based landslide identification product for nepal. *Earth Interactions* **23**(3), 1–21.

Firth, C. R. and Stewart, I. S. (2000) Postglacial tectonics of the Scottish glacio-isostatic uplift centre. *Quaternary Science Reviews* **19**(14-15), 1469–1493.

Fobert, M.-A., Singhroy, V. and Spray, J. G. (2021) InSAR Monitoring of Landslide Activity in Dominica. *Remote Sensing* **13**(4), 815.

Frattini, P. and Crosta, G. B. (2013) The role of material properties and landscape morphology on landslide size distributions. *Earth and Planetary Science Letters* **361**, 310–319.

Gabriel, E., Bonneu, F., Monestiez, P. and Chadœuf, J. (2016) Adapted kriging to predict the intensity of partially observed point process data. *Spatial Statistics* **18**, 54–71.

Galli, M., Ardizzone, F., Cardinali, M., Guzzetti, F. and Reichenbach, P. (2008) Comparing landslide inventory maps. *Geomorphology* **94**(3-4), 268–289.

Garnier, Simon, Ross, Noam, Rudis, Robert, Camargo, Pedro, A., Sciaini, Marco, Scherer and Cédric (2024) *viridis(Lite) - Colorblind-Friendly Color Maps for R.* viridis package version 0.6.5.

Gelman, A., Hwang, J. and Vehtari, A. (2014) Understanding predictive information criteria for Bayesian models. *Statistics and Computing* **24**(6), 997–1016.

Gibbens, S. (2019) This Caribbean island is on track to become the world's first 'hurricane-proof' country. *National Geographic* **Science**.

Gneiting, T., Balabdaoui, F. and Raftery, A. E. (2007) Probabilistic forecasts, calibration and sharpness. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **69**(2), 243–268.

Goetz, J., Brenning, A., Petschko, H. and Leopold, P. (2015) Evaluating machine learning and statistical prediction techniques for landslide susceptibility modeling. *Computers & Geosciences* **81**, 1–11.

Goldsmith, J., Scheipl, F., Huang, L., Wrobel, J., Di, C., Gellar, J., Harezlak, J., McLean, M. W., Swihart, B., Xiao, L., Crainiceanu, C., Reiss, P. T. and Cui, E. (2024) *refund: Regression with Functional Data.* R package version 0.1-37.

GRASS Development Team (2024) *Geographic Resources Analysis Support System (GRASS GIS) Software, Version 8.4.* Open Source Geospatial Foundation, USA.

Guttorp, P. and Gneiting, T. (2006a) Studies in the History of Probability and Statistics XLIX on the Matérn Correlation family. *Biometrika* **93**(4), 989–995.

Guttorp, P. and Gneiting, T. (2006b) Studies in the History of Probability and Statistics XLIX on the Matérn Correlation family. *Biometrika* **93**(4), 989–995.

Guzzetti, F., Carrara, A., Cardinali, M. and Reichenbach, P. (1999) Landslide hazard evaluation: a review of current techniques and their application in a multi-scale study, central Italy. *Geomorphology* **31**(1), 181–216.

Guzzetti, F., Gariano, S. L., Peruccacci, S., Brunetti, M. T., Marchesini, I., Rossi, M. and Melillo, M. (2020) Geographical landslide early warning systems. *Earth-Science Reviews* **200**, 102973.

Guzzetti, F., Mondini, A. C., Cardinali, M., Fiorucci, F., Santangelo, M. and Chang, K.-T. (2012) Landslide inventory maps: new tools for an old problem. *Earth-Science Reviews* **112**(1-2), 42–66.

Guzzetti, F., Reichenbach, P., Cardinali, M., Galli, M. and Ardizzone, F. (2005) Probabilistic landslide hazard assessment at the basin scale. *Geomorphology* **72**, 272–299.

Haskard, K. A., Cullis, B. R. and Verbyla, A. P. (2007) Anisotropic Matérn correlation and spatial prediction using REML. *Journal of Agricultural, Biological, and Environmental Statistics* **12**, 147–160.

He, J., Tanyas, H., Huang, D. and Lombardo, L. (2025) Space-time explainable modelling of regional hillslope deformation, an example from the tibetan plateau. *Remote Sensing of Environment* **329**, 114924.

Hijmans, R. J. (2024) *raster: Geographic Data Analysis and Modeling*. R package version 3.6-30.

Hosmer, D. W. (2000) Applied logistic regression, 2 .

Illian, J., Penttinen, A., Stoyan, H. and Stoyan, D. (2008) *Statistical Analysis and Modelling of Spatial Point Patterns*. John Wiley & Sons, Ltd.

Illian, J. B., Sørbye, S. H. and Rue, H. (2012) A toolbox for fitting complex spatial point process models using integrated nested Laplace approximation (INLA) .

Iverson, R. M. (1997) The physics of debris flows. *Reviews of Geophysics* **35**(3), 245–296.

Kirschbaum, D., Stanley, T. and Zhou, Y. (2015) Spatial and temporal analysis of a global landslide catalog. *Geomorphology* **249**, 4–15.

Klein, C. (2023) Natural disasters in Japan - statistics and facts.

Krainski, E., Gómez-Rubio, V., Bakka, H., Lenzi, A., Castro-Camilo, D., Simpson, D., Lindgren, F. and Rue, H. (2018) *Advanced spatial modeling with stochastic partial differential equations using R and INLA*. Chapman and Hall/CRC.

Krainski, E. T., Lindgren, F. and Rue, H. (2025) *INLAspacetime: Spatial and Spatio-Temporal Models using 'INLA'*. R package version 0.1.12.002.

Lari, S., Frattini, P. and Crosta, G. (2014) A probabilistic approach for landslide hazard analysis. *Engineering Geology* **182**, 3–14.

Lima, P., Steger, S. and Glade, T. (2021) Counteracting flawed landslide data in statistically based landslide susceptibility modelling for very large areas: a national-scale assessment for Austria. *Landslides* **18**(11), 3531–3546.

Lima, P., Steger, S., Glade, T., Tilch, N., Schwarz, L. and Kociu, A. (2017) Landslide susceptibility mapping at national scale: a first attempt for Austria. In *Advancing Culture of Living with Landslides: Volume 2 Advances in Landslide Science*, pp. 943–951.

Lin, Q., Lima, P., Steger, S., Glade, T., Jiang, T., Zhang, J., Liu, T. and Wang, Y. (2021) National-scale data-driven rainfall induced landslide susceptibility mapping for China by accounting for incomplete landslide data. *Geoscience Frontiers* **12**(6), 101248.

Lindgren, F., Bolin, D. and Rue, H. (2022) The SPDE approach for Gaussian and non-Gaussian fields: 10 years and still running. *Spatial Statistics* **50**, 100599.

Lindgren, F., Rue, H. and 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: Statistical Methodology* **73**(4), 423–498.

Liu, J. and Vanhatalo, J. (2020) Bayesian model based spatiotemporal survey designs and partially observed log Gaussian Cox process. *Spatial Statistics* **35**, 100392.

Liu, Z. and Rue, H. (2022) Leave-group-out cross-validation for latent Gaussian models. *arXiv preprint arXiv:2210.04482* .

Lombardo, L., Bakka, H., Tanyas, H., van Westen, C., Martin Mai, P. and Huser, R. (2019a) Geostatistical Modeling to Capture Seismic-Shaking Patterns From Earthquake-Induced Landslides. *Journal of Geophysical Research: Earth Surface* **124**, 1958–1980.

Lombardo, L., Cama, M., Maerker, M. and Rotigliano, E. (2014) A test of transferability for landslides susceptibility models under extreme climatic events: application to the Messina 2009 disaster. *Natural Hazards* **74**, 1951–1989.

Lombardo, L., Opitz, T., Ardizzone, F., Guzzetti, F. and Huser, R. (2020) Space-time landslide predictive modelling. *Earth-Science Reviews* **209**.

Lombardo, L., Opitz, T. and Huser, R. (2018) Point process-based modeling of multiple debris flow landslides using INLA: an application to the 2009 Messina disaster. *Stochastic Environmental Research and Risk Assessment* **32**, 2179–2198.

Lombardo, L., Opitz, T. and Huser, R. (2019b) Numerical Recipes for Landslide Spatial Prediction Using R-INLA: A Step-by-Step Tutorial. *Spatial Modeling in GIS and R for Earth and Environmental Sciences* pp. 55–83.

Lombardo, L., Tanyas, H., Huser, R., Guzzetti, F. and Castro-Camilo, D. (2021) Landslide size matters: A new data-driven, spatial prototype. *Engineering Geology* **293**.

Lv, L., Chen, T., Dou, J. and Plaza, A. (2022) A hybrid ensemble-based deep-learning framework for landslide susceptibility mapping. *International Journal of Applied Earth Observation and Geoinformation* **108**, 102713.

Malone, A. (1997) Risk management and slope safety in hong kong (keynote address). In *Proceedings of the HKIE Geotechnical Division Annual Seminar on Slope Engineering in Hong Kong*, pp. 305–310.

Martha, T. R., van Westen, C. J., Kerle, N., Jetten, V. and Kumar, K. V. (2013) Landslide hazard and risk assessment using semi-automatically created landslide inventories. *Geomorphology* **184**, 139–150.

Matérn, B. (1986a) Some remarks on the topographic variation. In *Spatial Variation*, pp. 51–68. Springer.

Matérn, B. (1986b) Spatial variation, volume 36 of lecture notes in statistics.

McCullagh, P. (2019) *Generalized linear models*. Routledge.

McLean, M. W., Hooker, G., Staicu, A.-M., Scheipl, F. and Ruppert, D. (2014) Functional generalized additive models. *Journal of Computational and Graphical Statistics* **23**(1), 249–269.

Meyer, R. (2014) Deviance information criterion (DIC). *Wiley StatsRef: Statistics Reference Online* pp. 1–6.

Miller, D. L., Glennie, R. and Seaton, A. E. (2020) Understanding the stochastic partial differential equation approach to smoothing. *Journal of Agricultural, Biological and Environmental Statistics* **25**(1), 1–16.

Møller, J., Syversveen, A. R. and Waagepetersen, R. P. (1998) Log Gaussian Cox processes. *Scandinavian Journal of Statistics* **25**(3), 451–482.

Montgomery, D. R. and Dietrich, W. E. (1994) A physically based model for the topographic control on shallow landsliding. *Water Resources Research* **30**(4), 1153–1171.

Moreno, M., Lombardo, L., Crespi, A., Zellner, P. J., Mair, V., Pittore, M., van Westen, C. and Steger, S. (2024) Space-time data-driven modeling of precipitation-induced shallow landslides in South Tyrol, Italy. *Science of the Total Environment* **912**, 169166.

Nava, L., Carraro, E., Reyes-Carmona, C., Puliero, S., Bhuyan, K., Rosi, A., Monserrat, O., Floris, M., Meena, S. R., Galve, J. P. *et al.* (2023) Landslide displacement forecasting using deep learning and monitoring data across selected sites. *Landslides* **20**(10), 2111–2129.

Neteler, M. and Mitasova, H. (2002) *Open source GIS: a GRASS GIS approach*. Volume 689. Springer Science & Business Media.

Neuland, H. (1976) A prediction model of landslips. *Catena* **3**(2), 215–230.

Oku, Y., Yoshino, J., Takemi, T. and Ishikawa, H. (2014) Assessment of heavy rainfall-induced disaster potential based on an ensemble simulation of Typhoon Talas (2011) with controlled track and intensity. *Natural Hazards and Earth System Sciences* **14**(10), 2699–2709.

Ozturk, U., Pittore, M., Behling, R., Roessner, S., Andreani, L. and Korup, O. (2021) How robust are landslide susceptibility estimates? *Landslides* **18**, 681–695.

O'Hagan, A. (2004) Bayesian statistics: principles and benefits. *Frontis* pp. 31–45.

Palamakumbura, R., Finlayson, A., Ciurean, R., Nedumpallile-Vasu, N., Freeborough, K. and Dashwood, C. (2021) Geological and geomorphological influences on a recent debris flow event in the ice-scoured mountain quaternary domain, Western Scotland. *Proceedings of the Geologists' Association* **132**(4), 456–468.

Pasch, R. J., Penny, A. B. and Berg, R. (2018) National hurricane center tropical cyclone report: Hurricane Maria. *National Oceanic And Atmospheric Administration and the National Weather Service* pp. 1–48.

Pebesma, E. and Bivand, R. S. (2023) *Spatial Data Science With Applications in R*. Chapman Hall.

Pebesma, E. J. *et al.* (2018) Simple features for R: standardized support for spatial vector data. *R J.* **10**(1), 439.

Pebesma, E. J. and Bivand, R. (2005) Classes and methods for spatial data in R. *R News* **5**(2), 9–13.

Pedersen, T. L. (2024) *patchwork: The Composer of Plots*. R package version 1.3.0.

Pennington, C., Freeborough, K., Dashwood, C., Dijkstra, T. and Lawrie, K. (2015) The national landslide database of Great Britain: Acquisition, communication and the role of social media. *Geomorphology* **249**, 44–51.

Pudasaini, S. P. and Krautblatter, M. (2021) The landslide velocity. *Earth Surface Dynamics Discussions* **2021**, 1–34.

Ramsay, J., Heckman, N. and Silverman, B. (1997) Spline smoothing with model-based penalties. *Behavior Research Methods, Instruments, & Computers* **29**(1), 99–106.

Reichenbach, P., Rossi, M., Malamud, B. D., Mihir, M. and Guzzetti, F. (2018) A review of statistically-based landslide susceptibility models. *Earth-Science Reviews* **180**, 60–91.

ReliefWeb (2018) Dominica: Landslides and floods triggered by Hurricane Maria (18 september 2017). Accessed: 16 May 2025.

Rue, H., Martino, S. and Chopin, N. (2009) Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **71**(2), 319–392.

Rue, H., Riebler, A., Sørbye, S. H., Illian, J. B., Simpson, D. P. and Lindgren, F. K. (2017) Bayesian computing with INLA: a review. *Annual Review of Statistics and its Application* **4**, 395–421.

Samia, J., Temme, A., Bregt, A., Wallinga, J., Guzzetti, F., Ardizzone, F. and Rossi, M. (2017) Characterisation and quantification of path dependency in landslide susceptibility. *Geomorphology* **292**, 16–24.

Sarkar, D. (2008) *Lattice: Multivariate Data Visualization with R*. New York: Springer. ISBN 978-0-387-75968-5.

Sarkar, D. and Andrews, F. (2022) *latticeExtra: Extra Graphical Utilities Based on Lattice*. R package version 0.6-30.

Scheip, C. and Wegmann, K. (2022) Insights on the growth and mobility of debris flows from repeat high-resolution lidar. *Landslides* **19**(6), 1297–1319.

Scheipl, F., Staicu, A.-M. and Greven, S. (2015) Functional additive mixed models. *Journal of Computational and Graphical Statistics* **24**(2), 477–501.

Shano, L., Raghuvanshi, T. K. and Meten, M. (2020) Landslide susceptibility evaluation and hazard zonation techniques–a review. *Geoenvironmental Disasters* **7**, 1–19.

Sigrist, F. (2022) Latent gaussian model boosting. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(2), 1894–1905.

Simpson, D., Illian, J. B., Lindgren, F., Sørbye, S. H. and Rue, H. (2016) Going off grid: computationally efficient inference for log-Gaussian Cox processes. *Biometrika* **103**(1), 49–70.

Simpson, D., Lindgren, F. and Rue, H. (2012) In order to make spatial statistics computationally feasible, we need to forget about the covariance function. *Environmetrics* **23**(1), 65–74.

Simpson, D., Rue, H., Riebler, A., Martins, T. G. and Sørbye, S. H. (2017) Penalising model component complexity: A principled, practical approach to constructing priors. *Statistical Science* pp. 1–28.

Sosio, R., Crosta, G. and Frattini, P. (2007) Field observations, rheological testing and numerical modelling of a debris-flow event. *Earth Surface Processes and Landforms* **32**(2), 290–306.

Steger, S., Brenning, A., Bell, R. and Glade, T. (2016) The propagation of inventory-based positional errors into statistical landslide susceptibility models. *Natural Hazards and Earth System Sciences* **16**(12), 2729–2745.

Steger, S., Mair, V., Kofler, C., Pittore, M., Zebisch, M. and Schneiderbauer, S. (2021) Correlation does not imply geomorphic causation in data-driven landslide susceptibility modelling–benefits of exploring landslide data collection effects. *Science of the Total Environment* **776**, 145935.

Stein, M. L. (2012) *Interpolation of spatial data: some theory for kriging*. Springer Science & Business Media.

Steyerberg, E. W., Eijkemans, M. J. and Habbema, J. D. F. (1999) Stepwise selection in small data sets: a simulation study of bias in logistic regression analysis. *Journal of Clinical Epidemiology* **52**(10), 935–942.

Sullivan, T. (2020) A brief introduction to weak formulations of pdes and the finite element method. *University of Warwick* .

Sun, Z., Yang, X., Zhao, Z. and Zhang, B. (2022) Study on deformation characteristics of different angle loess slopes under wet-dry alternation. *Arabian Journal of Geosciences* **15**(12), 1166.

Tang, H., Hu, X., Xu, C., Li, C., Yong, R. and Wang, L. (2014) A novel approach for determining landslide pushing force based on landslide-pile interactions. *Engineering Geology* **182**, 15–24.

Tanyaş, H., Hill, K., Mahoney, L., Fadel, I. and Lombardo, L. (2022) The world's second-largest, recorded landslide event: Lessons learnt from the landslides triggered during and after the 2018 Mw 7.5 Papua New Guinea earthquake. *Engineering Geology* **297**, 106504.

Tanyas, H. and Lombardo, L. (2020) Completeness index for earthquake-induced landslide inventories. *Engineering Geology* **264**, 105331.

Taylor, F. E., Malamud, B. D., Witt, A. and Guzzetti, F. (2018) Landslide shape, ellipticity and length-to-width ratios. *Earth Surface Processes and Landforms* **43**(15), 3164–3189.

Titti, G., van Westen, C., Borgatti, L., Pasuto, A. and Lombardo, L. (2021) When Enough Is Really Enough? On the Minimum Number of Landslides to Build Reliable Susceptibility Models. *Geosciences* **11**(11), 469.

Turel, M. and Frost, J. D. (2011) Delineation of slope profiles from digital elevation models for landslide hazard analysis. In *Geo-Risk 2011: Risk Assessment and Management*, pp. 829–836.

Van Niekerk, J., Krainski, E., Rustand, D. and Rue, H. (2023) A new avenue for Bayesian inference with INLA. *Computational Statistics & Data Analysis* **181**, 107692.

Van Westen, C. and Zhang, J. (2018) Tropical Cyclone Maria. *Inventory of Landslides and Flooded Areas UNITAR Map Product ID* **2762**.

Waagepetersen, R. (2004) Convergence of posteriors for discretized log Gaussian Cox processes. *Statistics & Probability Letters* **66**(3), 229–235.

Wickham, H. (2016) *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York. ISBN 978-3-319-24277-4.

Wickham, H., François, R., Henry, L., Müller, K. and Vaughan, D. (2023) *dplyr: A Grammar of Data Manipulation*. R package version 1.1.4.

Wickham, H., Vaughan, D. and Girlich, M. (2024) *tidyr: Tidy Messy Data*. R package version 1.3.1.

Wood, S. N. (2011) Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society (B)* **73**(1), 3–36.

Wood, S. N. (2017) *Generalized additive models: an introduction with R*. Chapman and Hall/CRC.

Xie, Y. (2024) *knitr: A General-Purpose Package for Dynamic Report Generation in R*. R package version 1.48.

Yoshimatsu, H. and Abe, S. (2006) A review of landslide hazards in Japan and assessment of their susceptibility using an analytical hierarchic process (ahp) method. *Landslides* **3**(2), 149–158.

Zhang, H. (2004) Inconsistent estimation and asymptotically equal interpolations in model-based geostatistics. *Journal of the American Statistical Association* **99**(465), 250–261.

Zou, K. H., Resnic, F. S., Talos, I.-F., Goldberg-Zimring, D., Bhagwat, J. G., Haker, S. J., Kikinis, R., Jolesz, F. A. and Ohno-Machado, L. (2005) A global goodness-of-fit test for receiver operating characteristic curve analysis via the bootstrap method. *Journal of Biomedical Informatics* **38**.