

HIGH-RESOLUTION BAYESIAN MAPPING OF LANDSLIDE HAZARD WITH UNOBSERVED TRIGGER EVENT

BY THOMAS OPITZ^{1,a} , HAAKON BAKKA^{2,b} , RAPHAËL HUSER^{3,c}  AND
LUIGI LOMBARDO^{4,d} 

¹*Biostatistics and Spatial Processes, INRAE, thomas.opitz@inrae.fr*

²*Department of Mathematics, University of Oslo, haakoncb@math.uio.no*

³*Computer, Electrical and Mathematical Sciences and Engineering (CEMSE) Division, King Abdullah University of Science and Technology (KAUST), raphael.huser@kaust.edu.sa*

⁴*Faculty of Geo-Information Science and Earth Observation, University of Twente, l.lombardo@utwente.nl*

Statistical models for landslide hazard enable mapping of risk factors and landslide occurrence intensity by using geomorphological covariates available at high spatial resolution. However, the spatial distribution of the triggering event (e.g., precipitation or earthquakes) is often not directly observed. In this paper we develop Bayesian spatial hierarchical models for point patterns of landslide occurrences using different types of log-Gaussian Cox processes. Starting from a competitive baseline model that captures the unobserved precipitation trigger through a spatial random effect at slope unit resolution, we explore novel complex model structures that take clusters of events arising at small spatial scales into account as well as nonlinear or spatially-varying covariate effects. For a 2009 event of around 5000 precipitation-triggered landslides in Sicily, Italy, we show how to fit our proposed models efficiently, using the integrated nested Laplace approximation (INLA), and rigorously compare the performance of our models both from a statistical and applied perspective. In this context we argue that model comparison should not be based on a single criterion and that different models of various complexity may provide insights into complementary aspects of the same applied problem. In our application our models are found to have mostly the same spatial predictive performance, implying that key to successful prediction is the inclusion of a slope-unit resolved random effect capturing the precipitation trigger. Interestingly, a parsimonious formulation of space-varying slope effects reflects a physical interpretation of the precipitation trigger: in subareas with weak trigger, the slope steepness is shown to be mostly irrelevant.

1. Introduction. Landslides are ubiquitous in mountainous regions around the globe. A landslide can be defined as a mass movement of rock, debris or earth down a slope under the influence of gravity (Varnes (1958)). It can occur whenever the friction along a sliding plane is less than the downward-facing forces acting on the same surface (Bout et al. (2018)). The term landslide includes several modes of mass movement, namely, falls, topples, slides, spreads or flows (Hungri, Leroueil and Picarelli (2014)). In this work we consider debris flows, and we develop spatial predictive models for a particularly severe event in 2009 during which multiple landslides were triggered simultaneously by heavy rainfall in a region of the island of Sicily, Italy. Debris flows are phenomena that consist of a mixed mass of water, soil and fragmented rocks which: (i) flows over mountainsides, (ii) funnels into channels, (iii) entrains objects while propagating downhill and (iv) lays waste through its path before (v) stopping over valley floors (van den Bout et al. (2021)). As debris flows are primarily triggered by precipitation, rain discharge is positively correlated with landslide frequency. The collection

Received June 2020; revised October 2021.

Key words and phrases. Integrated nested Laplace approximation (INLA), landslide, log-Gaussian Cox process, model complexity, space-varying regression, triggering event.

of debris flows that we study was caused by an extreme rainfall event with a return period estimated at 30 years (Cama et al. (2015)); see Section 2 for more details on the dataset.

The operational use of statistical models for natural hazards is to combine the predictive information on landslide occurrences (Reichenbach et al. (2018)) with the distribution of elements at risk, such as infrastructure and built-up areas, and/or with the distribution of human population (Corominas et al. (2014)) to estimate potential losses. Most of the current approaches to mapping landslide hazard exploit auxiliary information from geomorphological covariates and focus on one of its components—known as the landslide susceptibility—through the modeling of presence-absence information (Atkinson and Masari (1998), Ayalew and Yamagishi (2005), Castro Camilo et al. (2017), Goetz et al. (2015), Rossi et al. (2010)); see Reichenbach et al. (2018) for a recent review. Such approaches are predominantly based on machine learning techniques using binary classification over fine pixel grids, where subsampling of zero observations is often inevitable to cope with high spatial resolution and highly imbalanced designs. Recently, Lombardo, Opitz and Huser (2018, 2019) and Lombardo et al. (2019, 2020) introduced the “intensity” concept for spatial landslide prediction by focusing on event counts and not only presence-absence data. Specifically, they proposed a novel probabilistic approach based on Bayesian hierarchical models, where landslides are viewed as spatial or spatiotemporal point processes of log-Gaussian Cox type. Using the integrated nested Laplace approximation (Illian, Sørbye and Rue (2012), Rue, Martino and Chopin (2009)), they developed accurate statistical inference with high grid resolution and with sophisticated latent structures for capturing intensity variations not explained by observed covariates. In this paper we use one of their models as a highly competitive baseline and explore various more complex model extensions described below. While observed covariates are available at high pixel resolution, spatial random effects can be resolved at multiple scales. Here, we use *slope units* (SUs, Amato et al. (202), Carrara et al. (1995)), which allow fast computations by being at lower spatial resolution than pixels, while delimiting physically-motivated zones that are relevant to the landsliding process which is known to show a relatively homogeneous response to slope instabilities within each SU. SUs are commonly used in landslide science (and more generally geomorphology), because of the empirical evidence that landslides occur on slopes (Guzzetti and Reichenbach (1994)).

We study nontrivial model extensions with respect to two important aspects: first, we include spatially unstructured effects at pixel or SU scale to capture residual spatial clustering at small scales; second, we construct complex models with nonlinear or space-varying covariate effects in order to improve the baseline model’s predictive performance and allow for new insights and interpretations from an applied perspective. Space-varying regression has been established as a useful concept when the response to a covariate is not constant in space (Gamerman, Moreira and Rue (2003), Gelfand et al. (2003)). In our context we consider it as a natural—yet difficult to implement—solution to account for the spatially-varying influence of the landslide trigger (such as a heavy precipitation event), which is usually not observed at good spatial resolution in the study area, or not observed at all. In this paper we conduct a thorough and rigorous structural analysis based on models that explore how the unobserved trigger interacts with observed covariate effects, here highlighted with the example of the slope steepness. Throughout our statistical analysis our aim is to integrate physical understanding about the landsliding process into the model structure, for example, by assuming that the slope steepness should become irrelevant as a predictor in places where there is only a weak—or no—trigger influence. Such improvements have never been considered in the landslide modeling literature, despite being of high practical relevance. More generally, we also discuss diagnostics to comprehensively compare the goodness-of-fit and spatial predictive performance of a range of models of varying complexity in order to study improvements with respect to the baseline. We argue that decisions should not be based on a single criterion




FIG. 1. Satellite image of the study area on the island of Sicily, Italy.

but rather on a combined assessment of several criteria to more objectively appreciate the relative strengths and weaknesses of models. Different models may in fact give complementary insights into the statistical or physical behavior of landslide activations.

In the remainder of the paper, Section 2 describes the landslide data and geomorphological covariates that we use in our analysis. Then, Section 3 provides the necessary background theory on log-Gaussian Cox processes and presents our new models, while Section 4 discusses implementation and model comparison using the integrated nested Laplace approximation. In Section 5 we compare and interpret the fitted models, which we exploit to map the landslide intensity and slope-related risk factors. We conclude with some discussion in Section 6.

2. Landslides data and predictor variables.

2.1. Precipitation-triggered landslide occurrences. On October 1, 2009, a major rainfall discharge occurred in an area of around 100 km^2 on the island of Sicily (Southern Italy, see Figure 1 for a satellite image of the study area), with approximately 250 mm of rain measured at nearby weather stations. This weather event followed two relatively smaller precipitation events just one and two weeks before, with about 190 mm and 75 mm of rain, respectively (Lombardo et al. (2016a)). Within just a few hours this extreme precipitation event triggered several thousands of rapid shallow landslides that led to the death of 37 people and economic infrastructure damage of around half a billion Euro. Using remote sensing images before and after the event, the identification of 4874 separate debris-flow landslides (Hungri, Leroueil and Picarelli (2014)) was made possible. So-called landslide identification points (LIPs, Lombardo et al. (2014)) were then extracted from remotely sensed images (at 0.25m resolution) for each mass movement. Precisely, the triggering location was set to the point of highest altitude in the area affected by the movement. The left panel of Figure 2 shows a digital elevation model of the study area and the LIP inventory at 15 m pixel resolution. For the coordinate projection system we use the metric Monte Mario projection for the Zone 2 of Italy (EPSG 3004); we then transform coordinates to km units and shift the values of both coordinates such that they start at zero for our study area. Most LIPs were recorded in distinct pixels, but with a few exceptions: 353 pixels contained two landslides, 44 pixels contained three landslides and two pixels contained four landslides.




FIG. 2. Landslide inventory and geomorphological features. The maps show the digital elevation model (DEM, left) and the Slope Steepness (right) as well as the landslide inventory (with black dots representing the LIPs).

2.2. *Geomorphological covariate information.* We use covariate information that has been aggregated to a $15\text{ m} \times 15\text{ m}$ grid based on a digital elevation model (DEM) with 2 m resolution. This grid resolution has been shown to be sufficiently fine to avoid degrading the predictive performance of models in this study area (Arnone et al. (2016), Cama et al. (2016), Lombardo et al. (2016b)) while allowing for reasonably fast inference. Lombardo, Opitz and Huser (2018) provide a more detailed description of the calculation and meaning of covariates. Aggregation from two meters to 15 meters was done by averaging values for continuous covariates or by selecting the prevailing category for categorical covariates. Covariates are as follows, where we use upper-case notation throughout to refer to the names of these covariates: Elevation (or Digital Elevation Model, abbreviated DEM); Aspect, that is, the angle in $[0, 2\pi)$ describing the exposition of the area with respect to the north (Zevenbergen and Thorne (1987)); Slope Steepness (Zevenbergen and Thorne (1987)); Planar Curvature (Heerdegen and Beran (1982)), which is measured perpendicular to the steepest slope angle and characterizes the convergence and divergence of flow across the surface; Profile Curvature (Heerdegen and Beran (1982)), which indicates the direction of maximum slope; Topographic Wetness Index (TWI) (Beven and Kirkby (1979)), which quantifies topographic properties related to hydrological processes using slope and upstream contributing area as input; Stream Power Index (SPI) (Moore, Grayson and Ladson (1991)), which takes similar input as TWI and measures more specifically the erosive power of flowing water; Landform (with 10 categories, see Wilson and Gallant (2000)); the distance of each pixel to the closest tectonic fault line (Dist2Fault, in m); Normalized Difference Vegetation Index (NDVI) (Rouse Jr. et al. (1974)), which measures the “greenness” of a landscape and serves as a proxy for vegetation; Lithology, that is, soil type with 22 categories, where rare soil types with less than 500 occurrence pixels have been summarized in a single class “other;” Land Use (with 13 categories). The choice of a $15\text{ m} \times 15\text{ m}$ grid yields a representation of the study area through 449,038 pixels. When using continuous covariates for modeling purposes, we scale them to have empirical mean 0 and empirical variance 1. Additional information about the covariates can be found in Section 1 of the Supplementary Material (Opitz et al. (2022)).




FIG. 3. Illustration of slope units (SUs) and of their adjacency structure. Left: Number of adjacent SUs to each SU indicated by color. Right: Adjacency graph.

Figure 2 shows the spatial distribution of the Elevation (left panel) and the Slope Steepness (right panel) on their original scale. In this work we stress the importance of accurately capturing the influence of the Slope Steepness which has a major effect on landslide activations. Landslides are very unlikely to be triggered in flat areas; they are much more likely on steeper slopes, and they become unlikely again on very steep slopes, since movable material has already gone through erosion and during previous mass movements. Below, we seek to construct models that can capture this nonmonotonic and highly nonlinear influence of the Slope Steepness.

In our Bayesian hierarchical models, described in Section 3, we additionally exploit modeling units at an intermediate resolution (between 15 m pixels and the full study area), known as *slope units* (SUs), which can be automatically extracted from the DEM data (Alvioli et al. (2016), Lombardo, Opitz and Huser (2018)). Precisely, the SU partition defines a physically-motivated and moderately-sized spatial discretization, here used for capturing latent random effects, such as the influence of the spatially-varying precipitation event. These SUs can be viewed as relatively homogeneous mapping units with respect to geomorphological and geophysical features that are relevant to landslide activations. In our study area we have 3484 SUs which are displayed on the left panel of Figure 3. Some landslides triggered within the same SU may be due to a joint triggering mechanism, which may potentially lead to some residual stochastic dependence in the landslide occurrence process (conditional on the geomorphological structure and the precipitation trigger), but such spatial dependence is likely to be very weak when considering events arising in separate SUs. In other words, landslide data in different SUs can be safely assumed to be conditionally independent given fixed and random effects. In the next section we describe our proposed log-Gaussian Cox process models.

3. Bayesian hierarchical modeling of landslide point patterns.

3.1. *Log-Gaussian Cox processes.* Spatial point processes are stochastic models for the occurrence of events in space, when the event positions are random but obey certain density patterns and small-scale clustering or inhibition behavior. We refer to the bounded study area as \mathcal{S} , shown in Figure 2. Point processes characterize the joint probability distribution of the number of points $N(A_i)$ in subareas A_1, A_2, \dots , that is, the marginal distribution and dependence structure of the random count variables $N(A_i)$. An important characteristic of point processes is the intensity function $\lambda(s)$, which determines the expected number of points over any area $A \subset \mathcal{S}$, that is, $\mathbb{E}\{N(A)\} = \int_A \lambda(s) ds$. For Poisson processes, $N(A)$ has the Poisson distribution for any set A , and the occurrence of points is independent, given the (deterministic) intensity function λ .

Log-Gaussian Cox processes (LGCPs, Møller, Syversveen and Waagepetersen (1998)) are Poisson processes with a stochastic intensity function $\Lambda(s)$ given by a log-Gaussian process. Their doubly-stochastic structure allows capturing spatial clustering of points, due to unobserved or unavailable predictor variables. LGCPs are convenient for Bayesian hierarchical modeling, where the latent Gaussian process may encompass both fixed effects of observed covariates $z_j(s)$ and random effects.

Throughout, we use the notation $x(s)$ and \mathbf{x} for random effects by adding context-specific subscripts and superscripts. Precisely, we follow the convention that $x(s)$ corresponds to the value of a random effect evaluated at the location $s \in \mathcal{S}$, and we write \mathbf{x} for the vector of the finite number of latent variables (following a multivariate Gaussian distribution) that are used to represent this random effect. For example, a random effect x_{SU} , described at SU resolution, corresponds to a multivariate Gaussian random vector with 3484 components, one for each SU, while $x_{\text{SU}}(s)$ is the value that corresponds to the SU containing the location $s \in \mathcal{S}$.

The log-intensities of our models are structured as

$$\log \Lambda(s) = \beta_0 + \sum_{j=1}^J \beta_j z_j(s) + \sum_{k=1}^K x_k(s), \quad s \in \mathcal{S}.$$

The random effects $x_k(s)$ may directly depend on location s or only indirectly through a covariate $z_j(s)$ observed at s , for instance, if $x_k(s)$ is used to capture the potentially nonlinear influence of a covariate. The probability density function of an observed finite point pattern $\mathcal{X} = (\mathbf{X}_1, \dots, \mathbf{X}_N)^T$, composed of a random but finite number $N \geq 0$ of points $\mathbf{X}_i \in \mathcal{S}$ in the observation window \mathcal{S} , corresponds to the expectation

$$f_{\text{LGCP}}(\mathcal{X}) = \mathbb{E}_{\Lambda} \left[\exp\left(-\int_{\mathcal{S}} \Lambda(s) ds\right) \prod_{i=1}^N \Lambda(\mathbf{X}_i) \right],$$

using the convention that $\prod_{i=1}^N \Lambda(\mathbf{X}_i) = 1$ if $N = 0$. Closed-form expressions of this expectation are not available in general, but Bayesian inference techniques, such as implementations based on the integrated nested Laplace approximation (Rue, Martino and Chopin (2009)), as used here, have been developed to approximate it numerically. In our Bayesian framework the Gaussian processes used to construct the log-Gaussian intensity function in LGCPs can also be viewed as prior distributions for deterministic components of the intensity function of a Poisson process.

3.2. *Penalized complexity priors for hyperparameters.* Hyperparameters (e.g., standard deviations of random effects) are important to control the behavior of components of the latent Gaussian process $\log \Lambda(s)$. They must be set or estimated carefully since different values for a hyperparameter may correspond to strong differences in the structure of the latent

components. Throughout all considered models, we here make systematic use of the concept of penalized complexity (PC) priors (Simpson et al. (2017)) to define the priors for hyperparameters. The rationale of PC priors is to control model complexity by assigning priors that are *not noninformative* to avoid overfitting. With PC priors, complexity of a model or model component is measured through a certain distance from a simpler baseline. PC priors are constructed to satisfy a list of natural criteria set out by Simpson et al. (2017) for priors of complex hierarchical models. Specifically, a PC prior controls “what a hyperparameter does” (i.e., how it acts on the model) rather than controlling its numerical value, which may differ according to different choices of parameterizations. For instance, if $\sigma > 0$ is a standard deviation, then σ^2 is the corresponding variance, and $1/\sigma^2$ is the corresponding precision. Nevertheless, the three values σ , σ^2 and $1/\sigma^2$ express the same general behavior, and their priors should behave accordingly. This is naturally achieved through PC priors. Thanks to their intuitive use and their contribution toward the stability of Bayesian inference procedures in sophisticated models, PC priors are now recommended as the standard choice for use with the integrated nested Laplace approximation (INLA).

With landslide models we define a simple baseline model for each latent component, typically corresponding to its absence (i.e., to a standard deviation of 0) such that all the latent variables pertaining to the component are set to 0. PC priors are constructed by assuming a constant-rate penalty for the distance of a model component with respect to its baseline. This distance is a measure of complexity, and the idea of PC priors, developed by (Simpson et al. (2017)), consists of expressing it in terms of the Kullback–Leibler divergence. The user has to specify the rate parameter for the penalty and can do so in an indirect, intuitive way by providing a threshold u and a probability $\alpha \in (0, 1)$ such that $\Pr(\text{hyperparameter} > u) = \alpha$. For illustration, let us consider a precision parameter $\tau > 0$ in a Gaussian random vector having distribution $\mathcal{N}(\mathbf{0}, \tau^{-1} \mathbf{Q}^{-1})$ with fixed precision matrix \mathbf{Q} . All PC priors used in our landslide models are of this type and can be obtained as follows. We define the baseline model through $\tau_0 = \infty$. To derive the PC prior from the Kullback–Leibler divergence of $\mathcal{N}(\mathbf{0}, \tau^{-1} \mathbf{Q}^{-1})$ with respect to $\mathcal{N}(\mathbf{0}, \tau_0^{-1} \mathbf{Q}^{-1})$, (Simpson et al. (2017)) consider finite but very large τ_0 . The result is a prior distribution with penalty rate parameter $\lambda > 0$, which incorporates τ_0 and can be fixed by the user as described above. The PC prior for the precision τ corresponds to a Fréchet distribution with density

$$\pi(\tau) = \frac{\lambda}{2} \tau^{-3/2} \exp(-\lambda \tau^{-1/2}), \quad \tau > 0,$$

which, in turn, corresponds to an exponential distribution with rate λ for the standard deviation parameter $\sqrt{1/\tau}$.

3.3. *Models.* We now present the baseline model and our proposed model extensions.

3.3.1. *M0 (baseline): Fixed effects and spatial random effect.* Our baseline model, called M0, has structure similar to the best model found by Lombardo, Opitz and Huser (2018), although here with relevant improvements through the choice of prior distributions penalizing model complexity. In the model extensions presented subsequently, we ensure easy comparison and consistency with M0 by keeping the same prior distributions for components that are in common. Before giving details on components of the log-intensity of M0, we provide its full formula,

$$(M0) \quad \log \Lambda_0(\mathbf{s}) = \sum_{j=1}^{m_{\text{cont}}} \beta_j^{\text{cont}} z_j^{\text{cont}}(\mathbf{s}) + \sum_{j=1}^3 \sum_{\ell=1}^{\ell_j} \beta_{j,\ell}^{\text{cat}} z_{j,\ell}^{\text{cat}}(\mathbf{s}) + x_{\text{Aspect}}^{\text{CRW1}}(\mathbf{s}) + x_{\text{SU}}^{\text{CAR}}(\mathbf{s}).$$

The intercept and the continuous covariates, $z_j^{\text{cont}}(\mathbf{s})$, including the Slope Steepness but with the exception of the circular Aspect variable, appear as fixed effects with coefficients β_j^{cont} , $j = 1, \dots, m_{\text{cont}}$. To guide the estimation algorithm for faster convergence and to stabilize the fitted model, we fix moderately informative Gaussian prior distributions with precision 1 and mean 0 for fixed effects, except for the intercept where the mean is -2 . This leads to posterior distributions that are strongly concentrated around zero for those covariates that cannot improve posterior predictions. Due to the log-link function, a latent Gaussian variable with a given mean, but either high or low variance, contributes very differently to the predictions obtained after applying the inverse link function $\exp(\cdot)$. Recall that if $x \sim \mathcal{N}(\mu, \sigma^2)$, then $\exp(x)$ has mean $\exp(\mu + \sigma^2/2)$. Our prior choice, therefore, reduces the gap between the generally higher variance in the prior model and the lower variance in the posterior model.

PC priors are used for the precision of the independent and identically distributed (i.i.d.) effects of the three categorical, nonordinal covariates (Lithology, Landform, Land Use), $z_{j,\ell}^{\text{cat}}(\mathbf{s})$, which possess a substantial number of categories ($\ell_1 = 22$, $\ell_2 = 10$ and $\ell_3 = 13$, respectively). Priors for their coefficients $\beta_{j,\ell}$, $\ell = 1, \dots, \ell_j$, $j = 1, 2, 3$, are centered at 0, and we impose a sum-to-zero constraint on the coefficients of each of the three factors to ensure identifiability. The priors for the three precision parameters, denoted by τ_{LITH} , τ_{LF} and τ_{LU} , respectively, are relatively informative and are determined by the a priori specification of $\Pr(\sqrt{1/\tau} > 1) = 0.01$ for the corresponding standard deviations $\sqrt{1/\tau}$. As outlined in Section 3.2, the PC prior for the standard deviation is given by an exponential distribution, whose rate parameter, corresponding to the above specification, amounts to 4.6. Therefore, we let the data drive the posterior distribution away from 0 if a clear signal is present in the data. This allows us to shrink the model toward a parsimonious formulation and to avoid unstable inference, such as excessively high posterior variances or difficulties in calculating Laplace approximations.

The Aspect covariate, $x_{\text{Aspect}}^{\text{CRW1}}(\mathbf{s})$, reports an angle within the interval $[0, 2\pi)$, which we here discretize into 16 equidistant bins, each spanning $2\pi/16 = 22.5^\circ$ for a near-continuous treatment. The prior model of this random effect is a cyclic first-order random walk (CRW1) over the bins with a sum-to-zero constraint for identifiability. Writing $\mathbf{x}_{\text{Aspect}}^{\text{CRW1}} = (x_{A,0}, \dots, x_{A,15})^T$, we characterize its multivariate Gaussian prior through the following conditional specification (where $\tilde{i} = i$ modulo 16):

$$x_{A,i} \mid \{x_{A,\tilde{i}-1}, x_{A,\tilde{i}+1}\} \sim \mathcal{N}\left(\frac{x_{A,\tilde{i}-1} + x_{A,\tilde{i}+1}}{2}, \frac{1}{\tau_A \tau_{A,0}}\right), \quad i = 0, \dots, 15,$$

where the constraint $\sum_{i=0}^{15} x_{A,i} = 0$ is imposed to ensure identifiability and where $\tau_{A,0} > 0$ is a fixed scaling constant such that $1/\tau_A$ corresponds to the marginal variance of the variables $x_{A,i}$, $i = 0, \dots, 15$. The CRW1-structure makes sure that the estimated piecewise constant curve is “smooth” by borrowing strength between neighboring classes. We set an informative PC prior distribution for $\mathbf{x}_{\text{Aspect}}^{\text{CRW1}}$ by specifying that, a priori, $\Pr(\sqrt{1/\tau_A} > 1) = 0.01$.

For the latent spatial effect $x_{\text{SU}}^{\text{CAR}}(\mathbf{s})$ structured at the SU level, we implement Besag’s classical conditional autoregressive (CAR) model (Besag (1975), Rue and Held (2005)). Writing $\mathbf{x}_{\text{SU}}^{\text{CAR}} = (x_{\text{SU},1}^{\text{CAR}}, \dots, x_{\text{SU},3484}^{\text{CAR}})^T$, this model links the value $x_{\text{SU},i}^{\text{CAR}}$ in SU i to adjacent SUs, described by the index set $\text{NB}(i)$ of size $n_i \geq 1$, as

$$x_{\text{SU},i}^{\text{CAR}} \mid \{x_{\text{SU},j}^{\text{CAR}}\}_{j \in \text{NB}(i)} \sim \mathcal{N}\left(\frac{1}{n_i} \sum_{j \in \text{NB}(i)} x_{\text{SU},j}^{\text{CAR}}, \frac{1}{n_i \tau_{\text{SU}} \tau_{\text{SU},0}}\right), \quad i = 1, \dots, 3484,$$

where the constraint $\sum_{i=1}^{3484} x_{\text{SU},i}^{\text{CAR}} = 0$ is imposed for identifiability and where $\tau_{\text{SU},0} > 0$ is a fixed scaling constant such that $1/\tau_{\text{SU}}$ corresponds to the (generalized) marginal variance (Sørbye and Rue (2014)) of the variables $x_{\text{SU},i}^{\text{CAR}}$, $i = 1, \dots, 3484$; the generalized marginal

variance is the square of the geometric mean of the (nonstationary) standard deviations of the variables. We set a moderately informative PC prior distribution for $\mathbf{x}_{\text{SU}}^{\text{CAR}}$ by specifying that, a priori, $\Pr(\sqrt{1/\tau_{\text{SU}}} > 5) = 0.01$. The PC prior for the standard deviation $\sqrt{1/\tau_{\text{SU}}}$ is again given by an exponential distribution (recall Section 3.2), whose rate parameter corresponding to the above specification amounts to 0.92.

The number of adjacent SUs, $|\text{NB}(i)|$, varies moderately in our dataset, with 90% of values between 3 and 8 and 60% between 4 and 6, while the minimum is 1 and the maximum is 19; see Figure 3. In general, this spatial random effect captures SU-resolved effects that cannot be explained by other model components, in particular, by observed covariates. With the landslides' data the spatial effect will absorb the local intensity variation of the precipitation trigger, which is, at most, weakly correlated with some of the other, geomorphological covariates.

3.3.2. Model 1: Spatially unstructured effects. Overdispersion in count data refers to the situation where the variance is larger than the mean which stands in contrast to the Poisson distribution whose mean and variance are equal. Conceptually, our LGCP models are defined over continuous space and exclude multiple events at the same location $s \in \mathcal{S}$ such that (theoretically) the notion of overdispersion does not apply. However, overdispersion in the counts for spatial units, such as pixels or SUs, can still arise if our intensity model is misspecified and fails to pick up all sources of spatial variation in the data. Our models assume constant intensity within pixels, and the spatial random effect has coarser resolution at the SU level. Pixel resolution is very high in our case and approximates continuous space, with only a very small proportion of pixels counting multiple events, such that we will not explore nonstationary behavior of the point process intensity within pixels. However, we propose to explore models with spatially unstructured effects at the pixel level or the SU level, which are capable of capturing sharp differences in intensity between neighboring pixels or SUs, respectively. In our first model extension we, therefore, include i.i.d. Gaussian variables in the latent linear predictor, either by adding one variable to each pixel (Model 1a) or by adding one variable to each SU (Model 1b). We estimate the precision parameters, $\tau_{\text{grid}}^{\text{iid}}$ and $\tau_{\text{SU}}^{\text{iid}}$, of these pixelwise and SU-wise unstructured effects, respectively. Writing $i(s)$ and $\text{SU}(s)$ to denote the pixel i and slope unit containing location $s \in \mathcal{S}$, respectively, these models are given as

$$(M1a) \quad \log \Lambda_{1a}(s) = \log \Lambda_0(s) + x_{\text{grid}}^{\text{iid}}\{i(s)\}, \quad \mathbf{x}_{\text{grid}}^{\text{iid}} \sim \mathcal{N}_{n_{\text{grid}}}(\mathbf{0}, I_{n_{\text{grid}}}/\tau_{\text{grid}}^{\text{iid}}),$$

$$(M1b) \quad \log \Lambda_{1b}(s) = \log \Lambda_0(s) + x_{\text{SU}}^{\text{iid}}\{\text{SU}(s)\}, \quad \mathbf{x}_{\text{SU}}^{\text{iid}} \sim \mathcal{N}_{n_{\text{SU}}}(\mathbf{0}, I_{n_{\text{SU}}}/\tau_{\text{SU}}^{\text{iid}}),$$

where sum-to-zero constraints are imposed on $\mathbf{x}_{\text{SU}}^{\text{iid}}$ and $\mathbf{x}_{\text{grid}}^{\text{iid}}$, $\mathbf{0}$ denotes the zero vector, I_n is the n -by- n identity matrix and $n_{\text{grid}} = 449,038$, $n_{\text{SU}} = 3484$. For both effects the precision parameter τ ($= \tau_{\text{grid}}^{\text{iid}}$ or $\tau_{\text{SU}}^{\text{iid}}$) is endowed with an informative PC prior determined by $\Pr(\sqrt{1/\tau} > 1) = 0.01$.

3.3.3. Model 2: Nonlinear effect. In this second model extension we replace the linear Slope Steepness effect of the form “ $\beta_{\text{Slope}} \times \text{Slope}(s)$ ” (with $\text{Slope}(s)$ a known covariate) by a nonlinear random effect $\mathbf{x}_{\text{Slope}}^{\text{RW1}}$, defined through a first-order random walk prior using 10 equidistant classes to partition Slope Steepness values. Denote the log-intensity of the baseline model without the linear Slope Steepness effect by $\log \Lambda_{0,-\text{Slope}}(s)$. Here, we consider the modified model

$$(M2) \quad \log \Lambda_2(s) = \log \Lambda_{0,-\text{Slope}}(s) + x_{\text{Slope}}^{\text{RW1}}(s),$$

which can capture nonlinear and, in particular, nonmonotonic influence of the Slope Steepness covariate. We set the prior distribution of the precision parameter $\tau_{\text{Slope}}^{\text{RW1}} > 0$ of $\mathbf{x}_{\text{Slope}}^{\text{RW1}}$ by analogy with the Aspect effect in M0, that is, $\Pr(\sqrt{1/\tau_{\text{Slope}}^{\text{RW1}}} > 1) = 0.01$.

3.3.4. *Model 3: Space-varying regression (SVR)*. Another extension of our baseline model is possible by keeping a linear coefficient for Slope but allowing it to vary over space. This allows the model to capture local variations of the strength of the Slope Steepness effect, due to the precipitation trigger. We keep the global linear Slope Steepness coefficient and add a spatially-varying correction, defined at the SU level, in the following model:

$$(M3) \quad \log \Lambda_3(s) = \log \Lambda_0(s) + \text{Slope}(s) \times x_{\text{Slope}}^{\text{CAR}}(s).$$

The role of the additional space-varying regression component is to locally correct the global effect where necessary. By analogy with the latent spatial effect $x_{\text{SU}}^{\text{CAR}}$, the prior on $x_{\text{Slope}}^{\text{CAR}}$ corresponds to a Gaussian process with CAR structure and with its own precision parameter $\tau_{\text{Slope}}^{\text{CAR}} > 0$ for which we set an informative prior distribution, according to $\Pr(\sqrt{1/\tau_{\text{Slope}}^{\text{CAR}}} > 0.1) = 0.01$. This corresponds to assigning an exponential distribution with rate 46 as the prior of the standard deviation $\sqrt{1/\tau_{\text{Slope}}^{\text{CAR}}}$. The penalty rate is very high since the addition of the SVR-component makes the model considerably more complex.

3.3.5. *Model 4: Nonlinear effect and space-varying regression*. In this model we combine both the nonlinear Slope Steepness effect of M2 and the SVR-coefficient component of M3 into a single model, leading to the following structure:

$$(M4) \quad \log \Lambda_4(s) = \log \Lambda_{0,-\text{Slope}}(s) + x_{\text{Slope}}^{\text{RW1}}(s) + \text{Slope}(s) \times x_{\text{Slope}}^{\text{CAR}}(s),$$

where hyperparameter priors are fixed as above.

3.3.6. *Model 5: Parsimonious space-varying regression (P-SVR)*. Finally, we construct a model similar to M4 but which links the latent spatial effect $x_{\text{SU}}^{\text{CAR}}(s)$ and the SVR component. If the latent spatial effect acts as a proxy for the precipitation trigger, then its low values indicate a weak or absent trigger effect, and then the Slope Steepness value becomes irrelevant since no landslides occur, whatever the geomorphological conditions. In this case the SVR may locally neutralize (i.e., counteract) the globally estimated Slope Steepness effect. We here consider the following parsimonious model:

$$(M5) \quad \log \Lambda_5(s) = \log \Lambda_{0,-\text{Slope}}(s) + x_{\text{Slope}}^{\text{RW1}}(s) + \beta \times \text{Slope}(s) \times x_{\text{SU}}^{\text{CAR}}(s)$$

with the interaction coefficient $\beta \in \mathbb{R}$ to be estimated. Unlike the more complex model M4, this model features only one single CAR effect, $x_{\text{SU}}^{\text{CAR}}(s)$, instead of the two a priori independent effects, $x_{\text{SU}}^{\text{CAR}}(s)$ and $x_{\text{Slope}}^{\text{CAR}}(s)$, in model M4 such that we consider it as a parsimonious variant of space-varying regression. The prior for the parameter β is set to be moderately informative; it is Gaussian with mean 1 and precision 10.

4. Approximate Bayesian inference.

4.1. *The integrated nested Laplace approximation (INLA)*. INLA (Rue, Martino and Chopin (2009)) has found widespread interest in a wide range of applications (Krainski et al. (2018), Lombardo, Opitz and Huser (2018), Moraga (2019), Opitz et al. (2018)) thanks to its ability to provide fast and accurate posterior inference for the general class of latent Gaussian models, including log-Gaussian Cox processes (Tierney and Kadane (1986)). The R-INLA package (see <http://www.r-inla.org/>), in which the core statistical methodology is efficiently and conveniently implemented, privileges sparse matrix calculations in large dimensions through systematic use of Gauss–Markov conditional independence structures. With hierarchically structured models, including several components with different structures at the latent layer, INLA is typically faster and simpler to tune than simulation-based Markov

chain Monte Carlo (MCMC) methods (Illian, Sørbye and Rue (2012), Opitz (2017), Rue and Held (2005), Rue, Martino and Chopin (2009), Rue et al. (2016)). INLA is an approximate inference method, but the approximation quality is generally superior to MCMC-based approaches when using similar computation times (Teng, Nathoo and Johnson (2017), Vranckx, Neyens and Faes (2019)), especially in cases where achieving good mixing properties within MCMC is difficult. The INLA method can be applied to models where the observations are conditionally independent with respect to a latent multivariate Gaussian random vector which parameterizes the observation model. The distribution of the observations is of Poisson type in our case; in general, INLA imposes only very mild conditions on the distribution of observations. INLA can be applied to a very wide class of latent Gaussian models (Rue et al. (2016)). The dimension of the latent Gaussian vector can be large, especially in spatial modeling; for instance, in the SU-based spatial random effects we have one latent variable for each of the 3484 SUs. Therefore, it is important that matrix computations related to the Gaussian model can be performed efficiently. With INLA this is possible if Gauss–Markov structures are used, that is, if the precision matrices (inverse covariance matrices) are sparse. By using a spatial discretization of the study area combined with a graphical structure, such as the SU structure with graph edges defined through spatial adjacency of SUs in our data, it is possible to define Gauss–Markov models for spatial data. Here, we use the classical spatial conditionally autoregressive model of Besag (1975) to estimate spatial random effects. The structure of log-Gaussian Cox processes perfectly fits into the framework of INLA (Illian, Sørbye and Rue (2012)), since the observed points of the point pattern (and the resulting Poisson-distributed counts over mapping units, such as pixels in our data) are conditionally independent with respect to the latent Gaussian process representing the logarithm of the intensity function. A detailed mathematical description of the estimation mechanism of INLA, which heavily draws on Laplace approximations (Tierney and Kadane (1986)), is given in Section 2 of the Supplementary Material (Opitz et al. (2022)).

We here write $\Lambda_i = \exp(\eta_i)$ for the stochastic point process intensity at pixel $i \in \{1, \dots, n_{\text{grid}}\}$ with $n_{\text{grid}} = 449,038$, where $\boldsymbol{\eta} = (\eta_1, \dots, \eta_{n_{\text{grid}}})^T$ denotes the vector of the pixel-based latent Gaussian log-intensities for the whole study area. We further use the notation $\mathbf{x}_{\text{full}} = (\boldsymbol{\eta}^T, \dots)^T$ for a vector with n_{full} components, where “...” refers to the variables of the additive random effects included in the log-intensity model, for example, $\mathbf{x}_{\text{SU}}^{\text{CAR}}$, and so forth. The vector of hyperparameters (i.e., precisions of CAR, RW1 and i.i.d. components) is denoted by $\boldsymbol{\theta}$ and has n_{θ} components. The distribution of pixel-based landslide counts y_i , collected into a vector $\mathbf{y} = (y_1, \dots, y_{n_{\text{grid}}})^T$, is assumed to be conditionally independent given the (random) intensity values, that is,

$$y_i \mid \Lambda_i \stackrel{\text{ind.}}{\sim} \text{Poisson}(C \Lambda_i), \quad i = 1, \dots, n_{\text{grid}},$$

where $C = (15 \text{ m})^2$ is a scaling factor corresponding to the area of one pixel. The principal inference goal is the calculation of the posterior densities of hyperparameters and of the components of \mathbf{x}_{full} , the latent vector with multivariate Gaussian prior distribution, that is,

$$(4.1) \quad \pi(\boldsymbol{\theta}_j \mid \mathbf{y}) = \int \pi(\mathbf{x}_{\text{full}}, \boldsymbol{\theta} \mid \mathbf{y}) \, d\mathbf{x}_{\text{full}} \, d\boldsymbol{\theta}_{-j}, \quad j = 1, \dots, n_{\theta},$$

$$(4.2) \quad \pi(x_i \mid \mathbf{y}) = \int \pi(x_i \mid \boldsymbol{\theta}, \mathbf{y}) \pi(\boldsymbol{\theta} \mid \mathbf{y}) \, d\boldsymbol{\theta}, \quad i = 1, \dots, n_{\text{full}}.$$

However, calculation of (4.1) and (4.2) is hampered by the high-dimensional numerical integration over the space $\mathbb{R}^{n_{\text{full}}}$ spanned by the Gaussian vector \mathbf{x}_{full} . Instead, INLA uses the Laplace approximation which corresponds to replacing integrand functions by suitable Gaussian density approximations. On the other hand, the integration with respect to the components of $\boldsymbol{\theta}$ is done through numerical integration schemes such that only a small number of hyperparameters can be estimated.

4.2. *Model comparison and selection.* First, we propose to compare models through the classical information criteria DIC and WAIC. These goodness-of-fit criteria take the effective dimension of the latent model into account, thus penalizing model complexity. Their close relationship to the predictive performance, measured through leave-one-out cross-validation, has been established, and WAIC is known to better take the stochasticity of the posterior predictive distributions into account (Gelman, Hwang and Vehtari (2014)). With INLA these quantities are calculated through sensible approximation techniques (Rue, Martino and Chopin (2009)).

To focus more directly on criteria evaluating the spatial predictive performance, we also devise a 10-fold cross-validation scheme. Specifically, we randomly partition the SUs into 10 folds, each containing (approximately) the same number of SUs. We calculate predictive scores for two mapping units, pixels and SUs, for the latter by aggregating observed and predicted counts over the pixels of each SU. At the pixel level we consider predictive scores that use either the predicted counts $\hat{\lambda}_i = \mathbb{E}(\Lambda_i | \mathbf{y})$, the predicted probabilities of within-pixel landslide occurrences $\hat{p}_i = 1 - \exp(-\hat{\lambda}_i)$ or the full INLA-based posterior predictive distribution of $\Lambda_i | \mathbf{y}$. Similarly, at the SU level we consider the predicted counts estimated as $\hat{\lambda}_{\text{SU}} = \sum_{i \in \text{SU}} \hat{\lambda}_i$, the predicted probabilities of within-SU landslide occurrences $\hat{p}_{\text{SU}} = 1 - \exp(-\hat{\lambda}_{\text{SU}})$ and the INLA-based posterior predictive distributions of $(\sum_{i \in \text{SU}} \Lambda_i) | \mathbf{y}$. An alternative approach for predictive diagnostics, studied by Leininger and Gelfand (2017), would be to construct hold-out sets by removing points at random from the point pattern; this is known as *thinning*. Here, we prefer the more challenging task of predicting entire spatially-contiguous areas where all data within SUs have been removed which is also more suited to assessing slopewise landslide hazard.

The INLA-based posterior predictive distributions are obtained by generating a large number of posterior samples of counts for each cross-validation fit. While INLA does not directly provide posterior samples because of its use of analytical and not simulation-based approximations, these can be generated conveniently (Rue et al. (2016)). Using R-INLA's internal, discrete approximations for posterior distributions of hyperparameters and latent Gaussian fields, the simulation algorithm first generates a realization of the hyperparameter vector; next, conditional on these hyperparameters, a latent Gaussian field is sampled; finally, counts are simulated from the pixel-based Poisson distributions with intensities defined according to the simulated latent Gaussian field. In what follows, cross-validation results, using simulations of the posterior predictive distributions, are based on 5000 samples of the full posterior model.

We consider four types of cross-validated predictive scores: the area-under-the-curve (AUC) (Fawcett (2006)) to measure prediction quality for the presence or absence of landslides within mapping units, the residual sum of squared errors (RSS) and the residual sum of absolute errors (RSA), both using predicted and observed counts and the continuous ranked probability score (CRPS, Gneiting and Katzfuss (2014)) using the predictive distribution functions and observed counts. The formulas for pixel-based RSS and RSA are as follows:

$$\text{RSS}_{\text{grid}} = \sum_{i=1}^{n_{\text{grid}}} (y_i - \hat{\lambda}_i)^2, \quad \text{RSA}_{\text{grid}} = \sum_{i=1}^{n_{\text{grid}}} |y_i - \hat{\lambda}_i|,$$

where $\hat{\lambda}_i = \int_0^\infty \pi_i \{\log(\lambda) | \mathbf{y}\} d\lambda$ with $\pi_i(\cdot | \mathbf{y})$ the INLA-based posterior density of $x_i = \log(\Lambda_i)$. The general CRPS formula for a single observation y_{obs} and a corresponding (posterior) predictive distribution $\hat{F}(y)$ from a model may be expressed as $\int_{-\infty}^\infty \{\hat{F}(y) - \mathbf{1}(y \geq y_{\text{obs}})\}^2 dy$. For the pixel-based CRPS in our case, we add up the CRPS values over all pixels and, therefore, use

$$\text{CRPS}_{\text{grid}} = \sum_{i=1}^{n_{\text{grid}}} \sum_{y=0}^\infty \left[\sum_{k=0}^y \int_0^\infty \exp(-\lambda) \frac{\lambda^k}{k!} \pi_i \{\log(\lambda) | \mathbf{y}\} \frac{1}{\lambda} d\lambda - \mathbf{1}(y \geq y_i) \right]^2.$$

Analogous formulas are used for SU-based criteria, where pixel-based observed counts y_i and intensities λ_i must be aggregated over the pixels for each SU. This requires resorting to the INLA-based joint posterior distributions of all x_i corresponding to the pixels in a given SU. Since such CRPS formulas are difficult to calculate analytically, we use posterior sampling as implemented in R-INLA and compute a Monte-Carlo approximation of CRPS values based on a large number of posterior samples (5000 as above).

The AUC considers only presence-absence data, which is a strong simplification for assessing the prediction of landslide counts, especially at the SU level where counts larger than one are frequent. By contrast, the other three measures rely on counts: while RSS and RSA focus on point predictions defined through the posterior mean of intensities at pixel level, CRPS also accounts for the uncertainty of the predictive distributions and yields good scores for models that provide predictions that are both calibrated (i.e., correct on average) and sharp (i.e., having little prediction uncertainty). We calculate these four predictive scores for each of the 10 folds, both at pixel and SU levels and, finally, we average the predictive scores of the 10 folds together. Lower final values correspond to better predictive performances.

5. Results and discussion.

5.1. *Computational requirements and general prior considerations.* INLA-based estimation of all models should be done on a computer with sufficiently high memory resources (at least 64 Gb). In our implementation we used four cores in parallel when estimating each model, and the resulting running times were between one and a half hour for the baseline model and around 18 hours for the model M1a with pixel-based i.i.d. random effect; the parsimonious-SVR model M5 took less than three hours. While estimation was possible in our setting without removing data from the sample, we note that estimation would require less resources, especially much less memory, if we performed subsampling of the pixels without observed landslides. If subsampling is performed in a stratified way to keep the loss of information small, then the original sample size could be divided by a factor of up to 10; see the discussion in Koh et al. (2021) for more background and references on weighted subsampling.

Based on the idea of penalizing model complexity in very complex models such as ours, we have opted for moderately informative priors in most cases. Specifically, model components that make models substantially more complex (e.g., the SVR-component in M4) have been endowed with informative priors. Since prior choice always involves some subjectivity, we have conducted several simulation experiments to assess how less informative priors would change results; their outcomes confirm that only slight differences would arise; see Section 3.2 of the Supplementary Material. Moreover, we have performed simulation experiments to assess if all the model components are appropriately identifiable in complex models, such as M4 and M5, and we give an affirmative answer; see Section 3.1 of the Supplementary Material (Opitz et al. (2022)).

5.2. *Comparison of models using basic diagnostics.* Table 1 reports results for our models fitted to the full dataset, including the precision parameter (i.e., inverse variance) of the estimated latent spatial effect (LSE) $\mathbf{x}_{\text{SU}}^{\text{CAR}}$ included in all models, the effective number of parameters (i.e., the effective dimension of the linear predictor when accounting for the dependence between the latent variables) and the two information criteria DIC and WAIC. The similar LSE precisions indicate that the variability of the LSE is relatively stable over different models, even with the most complex models. The effective number of parameters is relatively similar for all models, except the model with an i.i.d. effect resolved at the pixel

TABLE 1

Comparison of fitted models in terms of the estimated precision of the latent spatial effect (LSE) $\mathbf{x}_{\text{SU}}^{\text{CAR}}$ and information criteria (DIC and WAIC, respectively). The column n_{eff} denotes the effective number of parameters of the fitted model

Model	LSE precision	n_{eff}	DIC	WAIC
M0 (baseline)	0.35 (0.06)	1006	40,142	39,869
M1a (pixel-iid)	0.3 (0.02)	4166	41,925	40,834
M1b (SU-iid)	0.32 (0.02)	987	40,120	39,867
M2 (RW1)	0.31 (0.02)	1002	40,132	39,864
M3 (SVR)	0.31 (0.02)	1144	40,005	39,690
M4 (RW1-SVR)	0.31 (0.01)	1203	39,949	39,745
M5 (P-SVR)	0.34 (0.02)	1016	40,087	39,817

scale; recall that there is a large number 449,038 of pixels. DIC and WAIC values are relatively similar overall, although both information criteria consistently give preference to models with space-varying regression components. While DIC ranks first the most complex model M4 with independently specified RW1- and SVR-components for the Slope Steepness, WAIC prefers model M3 which includes only a fixed (i.e., global and linear) Slope Steepness effect and the SVR component.

We now also report and discuss estimated precision parameters for the specific components added in the models extending the baseline. In M1a with pixel-resolved i.i.d. effect, we estimate a posterior precision of about 1.5 for the i.i.d. component, which indicates the presence of a rather strong independent effect at the pixel scale, not explained by the aggregated view based on SUs, and without dependence spanning over neighboring pixels. By contrast, the precision of the SU-resolved i.i.d. effect in M1b is very high (180), indicating a relatively small contribution of this effect to the model. In M2 the precision of the nonlinear RW1-effect of Slope Steepness is 4.4. In M3 with an SVR component, the precision of the space-varying coefficient is 4.3. By jointly including the RW1- and SVR-effects of Slope Steepness in M4, we get RW1-precision of 5.8 and SVR precision of 2.4. The former is higher than without the SVR component (M2), indicating that the influence of Slope Steepness is now partially captured by the additional SVR component, whose precision is relatively low. We conjecture that this low precision shows that the SVR component captures the influence of Slope Steepness more easily than the global RW1-effect of Slope Steepness; moreover, the noncontinuous specification of the RW1-curve may require stronger variation of the SVR-component's contribution at relatively small spatial scales to smooth the RW1-effect. Finally, model M5 with a parsimonious SVR component has a RW1-precision estimated of 7.9, and the posterior mean of the β -coefficient is given by 0.13 with credible interval [0.10, 0.16]. Therefore, a significant transfer of predictive information has taken place from the LSE (whose precision is similar to the other models; recall Table 1) to the space-varying Slope Steepness influence, while the RW1-contribution has been reduced, compared to the other models with RW1-component. The parsimonious constraint linking the SVR to the latent spatial effect in M5 leads to an improved goodness-of-fit compared to the baseline M0 and the models with additional i.i.d. components, but, based on its DIC and WAIC values, we conclude that it cannot fully attain the high flexibility of models M3 and M4.

5.3. Influence of fixed effects. In Figure 4 we compare the estimated coefficients for the eight predictor variables included as fixed effects in our different models. The contribution of these covariates, and the associated uncertainty of their coefficients, are estimated to be very similar across the seven models studied here, except for the Slope Steepness, due to the major




FIG. 4. Estimated fixed effect coefficients (except the intercept) for the seven models (M0, M1a, M1b, M2, M3, M4, M5, from left to right in each panel corresponding to one of the covariates). Black dots show INLA-based posterior means, while vertical segments and blue endpoints indicate the size of 95% credible intervals. Fixed effect coefficients for Slope Steepness are fixed to 0 for some models and appear only through a black dot at level 0 in these cases.

differences in model structure with respect to the contribution of this covariate. Interestingly, the parsimonious SVR structure seems to have fully absorbed the influence of DEM into the Slope-SVR part of the model, $Slope(s) \times x_{Slope}^{CAR}(s)$. We do not show results for the categorical covariates, but the conclusions remain qualitatively similar.

For the models M2, M4 and M5 with a nonlinear Slope Steepness effect modeled through a RW1 component, Figure 5 shows the resulting estimated curves—for better readability of the plots, piecewise constant curves are replaced by piecewise continuous interpolations. Nonlinear influence is obvious from these plots and displays a similar bell shape in all three models, with intermediately steep slopes between 30 and 60 degrees presenting high relative risk of landslide occurrence. Models M4 and M5 include additional SVR-components to capture the slope-specific influence at SU resolution such that the estimated RW1 curves appear to be flatter which implies smaller variations in relative risk.

5.4. Spatial predictive performance comparison. Table 2 reports spatial predictive scores based on a 10-fold stratified cross-validation, where folds are composed of a random selection of entire SUs. For all models, pixel- and SU-based AUC values are very close to each other and reach around 0.9, indicating a very good performance for predicting the presence of




FIG. 5. Estimated RW1-effects for the Slope Steepness (from left to right, for models M2, M4 and M5). INLA-based posterior means are shown in black, 95% pointwise credible intervals in blue.

TABLE 2

Cross-validation-based comparison of spatial predictive performance of fitted models, with the score of the best-performing models shown in bold face. Scores are given with four significant digits. Mathematical details on the different scores are given in Section 4.2. Pixel-based and SU-based scores are denoted with the subscripts grid and SU, respectively

Model	AUC _{grid}	AUC _{SU}	RSA _{grid}	RSA _{SU}	RSS _{grid}	RSS _{SU}	CRPS _{grid}	CRPS _{SU}
M0 (baseline)	0.8958	0.9308	420.0	420.5	2614	2608	466.3	240.6
M1a (pixel-iid)	0.8960	0.9308	590.1	590.0	5628	5602	474.7	279.7
M1b (SU-iid)	0.8958	0.9309	420.6	421.2	2615	2635	466.3	240.5
M2 (RW1)	0.8956	0.9310	411.3	411.5	2481	2483	464.7	238.5
M3 (SVR)	0.8963	0.9305	441.1	441.7	2922	2936	467.9	242.8
M4 (RW1-SVR)	0.8964	0.9302	436.6	437	2770	2801	465.9	241.1
M5 (P-SVR)	0.8966	0.9311	430.1	429.8	2833	2819	466.4	241.3

landslides, especially at the SU level. When considering sums of absolute and squared errors (RSA and RSS, respectively), stronger differences between models arise, with model M2 with nonlinear Slope influence obtaining the best scores, and M1a with pixel-based i.i.d. effects performing substantially worse than the other models. Throughout, model M2 has a very good performance and has the best score except for the SU-based AUC value, although differences are rather small. Therefore, the inclusion of a nonlinear effect of the Slope Steepness, here implemented through a random effect with RW1 prior, is important for good prediction. The baseline model shows stable and good performance throughout and does not suffer from some relatively bad count-based scores arising for some of the extended models (except M2). Overall, the ranking of models, based on their predictive performance, looks quite different from the one based on goodness-of-fit measured through information criteria in Table 1. We emphasize that many of the differences in scores across models in Table 2 are not significant when we take into account the uncertainty, for instance, measured through the variability among the 10 scores. In particular, AUC-values show stronger variability across the hold-out sets for fixed models than across models. A possible reason is that very high stochasticity and complexity of prior models may lead to more unstable, noisy predictions. We recommend a careful inspection of the fitted models based on several criteria, for goodness-of-fit and for out-of-sample prediction. The “best” model M2 is more complex than the baseline model M0, but we add only a relatively small number of 10 latent components to achieve a nonlinear contribution of Slope Steepness. We also stress that, for our landslides data, the inclusion of i.i.d. effects (pixel- or SU-based) could not provide substantial improvements of goodness-of-fit or predictive performance. Finally, the models M3, M4 and M5, which possess extra flexibility thanks to a space-varying regression component, are relatively competitive overall, despite their relatively worse performance on RSS measures. Such models can still be useful by offering insights into the “physical” interaction of Slope Steepness with the unobserved precipitation trigger, as further explained in Section 5.6 below.

5.5. Landslide intensity mapping. In Figure 6, we show the INLA-based posterior mean of the estimated log-intensity (at pixel scale) of model M0 and the difference in log-intensity between the most complex model Model M4 (RW1-SVR) and M0. In the log-intensity of M0, the influence of geological structures, such as river valleys (very low intensity) and mountain ridges, comes out clearly. The spatial structure of the predicted values is dominated by the spatial effect $\chi_{\text{SU}}^{\text{CAR}}(s)$, capturing the influence of the spatial variation of the precipitation trigger. In this model the amplitude of log-predicted intensities is close to that of the posterior mean of the spatial effect which is evidence that the spatial effect is crucially necessary to




FIG. 6. INLA-based posterior mean of the estimated log-intensity at the pixel scale. Left: Baseline model M0. Right: Difference of log-intensity between models M4 (RW1-SVR) and M0. For better visualization a small number of difference values outside the interval $[-2, 2]$ have been replaced by -2 (if value < -2) or by 2 (if value > 2).

account for unobserved covariate effects and to locally counteract the influence of observed covariates. The latter may be locally misestimated, depending on the force of the precipitation trigger. Differences between models M0 and M4 are usually relatively minor, but in some small subareas, especially river valleys, model M4 has substantially higher log-intensity values.

5.6. Interpretation of the slope steepness contribution. The maps in Figure 7 show how the observed Slope Steepness variable contributes to the linear predictor, according to the different model structures (M0, M2, M3, M5). The SVR models, M3 and M5, lead to an overall contribution of Slope Steepness that is strongly conditioned on the precipitation trigger; weakly impacted areas such as the northwestern part of the study area have a very weak contribution close to 0. In contrast, the simpler models, M0 and M2, including a fixed effect or a RW1-effect, show spatial variation that is closer to the one of the Slope Steepness.

From a physical perspective we expect to observe, at most, a few landslides in areas with low or absent precipitation trigger, and, intuitively, the effect of covariates may, therefore, become irrelevant in such areas. While basic models, such as M0 and M2, are unable to capture this behavior, more complex models with space-varying covariate coefficients can adequately reflect such natural physical constraints and help to better highlight areas where a covariate substantially increases landslide hazard. Our model M5 (parsimonious SVR) is even *designed* to explicitly integrate this interaction between the precipitation trigger (represented here through the latent spatial effect (LSE), $x_{SU}^{CAR}(s)$) and the space-varying coefficient. The estimated coefficient $\hat{\beta} = 0.13$ in model M5 measures the strength of this interaction; its credible interval of $(0.10, 0.16)$ is relatively narrow and indicates substantial interaction.

At sites where the precipitation trigger is present, a strong response to Slope Steepness is expected for intermediate angles. In our models the LSE acts as a proxy for the influence of the precipitation trigger, and we, therefore, expect a strong interaction between the precipitation trigger and the Slope Steepness effect in the log-intensity of the point process. In Model




FIG. 7. Contribution of the Slope Steepness covariate to the linear predictor (here expressed through the posterior mean) according to the following models: M0 (fixed effect; upper left); M2 (RW1 effect; upper right); M3 (SVR; bottom left); M5 (P-SVR; bottom right). For better visualization, very small number of values outside the interval $[-2, 2]$ have been replaced by -2 (if value < -2) or by 2 (if value > 2).

M4 the LSE and the space-varying Slope Steepness effect are components without any prior dependence, but we can investigate the posterior model structure and, in particular, posterior correlation between these components. In Model M5 the interaction structure is fixed to a linear rescaling determined by the β -hyperparameter. The plots on the left-hand side of Figure 8 show how the space-varying Slope Steepness coefficient varies with respect to two other quantities: the SU-based estimated posterior mean of the LSE and the Slope Steepness value (which we average over the pixels of each SU to obtain SU-based values). In this figure the estimated Slope Steepness coefficients have been smoothed with a 2D kernel density estimator to produce a continuous color map. The original, nonsmoothed values were obtained at the points shown as small grey dots. A clear interaction pattern between the LSE and the SVR coefficient arises in M4, with lower LSE corresponding to lower SVR coefficient. We underline that the parsimonious model M5 is able to reproduce a very similar structure. The strong similarity of the results for the two models M4 and M5, despite M5 offering much less flexibility due to its rigid link between SVR-coefficient and LSE, persists in the plots on the right-hand side of Figure 8. They show the actual space-varying Slope Steepness effect (i.e.,




FIG. 8. Illustration of the interaction structure of Slope Steepness and the precipitation trigger (represented through the LSE) at the SU scale, with Slope Steepness values averaged for each SU. Left column: M4 (SVR). Right column: M5 (P-SVR). First row: Smoothed map of SVR coefficient $x_{\text{Slope}}^{\text{CAR}}(s)$. Second row: Smoothed map of $\text{Slope}(s) \times x_{\text{Slope}}^{\text{CAR}}(s)$. Small gray dots indicate the LSE values of the 3484 SUs.

$\text{Slope}(s) \times x_{\text{Slope}}^{\text{CAR}}(s)$ which corrects the global RW1-Slope Steepness effect. As expected, the plots show a correction toward lower relative landslide hazard for the combinations of high LSE/flat slopes and small LSE/steep slopes. A relatively strong correction toward higher relative landslide hazard is necessary for high LSE/intermediate slopes and relatively low LSE but Slope Steepness close to the pivotal value of 20 degrees (or slightly above). We conclude that the parsimonious model M5 offers a correction of the trigger-independent RW1-Slope Steepness effect that reflects known physical behavior and the typical landsliding response behavior to Slope Steepness, according to the (approximate) intervals 0–20 degrees, 20–35 degrees and > 35 degrees. Field and empirical evidence suggests that the first interval 0–20 degrees corresponds to slopes that are too flat for landslides to occur, even in the presence of a strong trigger event, while the last interval of > 35 degrees corresponds to slopes where, typically, material that is susceptible to sliding has already gone in the past. Therefore, the trigger primarily acts on slopes in the interval of 20–35 degrees where landslides occur most easily.

6. Conclusion. For georeferenced landslide events, we have proposed a Bayesian hierarchical point-process-based modeling framework to assess landslide hazard at high spatial resolution. Although substantial geomorphological covariate information is available in our case, it remains crucial for spatial prediction and for the interpretation of covariate effects to capture the latent activation pattern induced by the unobserved precipitation trigger. The framework of space-varying regression is appealing and useful, but care is needed to avoid overly complex models for which model components may be difficult to identify from the data and to interpret intuitively. When a spatial random effect captures the trigger intensity, it is natural to include it as a complement of space-varying regression such that local covariate effects can be locally removed when the trigger is absent or weak.

The inclusion of additional latent random effects without spatial dependence in the log-intensity allows us to model small-scale variability, either at the pixel level (here available at a high 15 m resolution) or at the coarser slope unit (SU) level. This enables capturing overdispersion that would otherwise lead to underestimated local variance of counts arising in the Poisson regression model used to spatially discretize the log-Gaussian Cox process. With our dataset, the inclusion of such i.i.d. effects did not substantially improve the baseline model. Rather, information criteria and some of the predictive performance measures indicated that model M1a with pixel-based i.i.d. effect was much worse which may be explained by the large number of additional independent latent variables. This leads to a strong increase in computational complexity of estimation and may hamper the identification of intermediate-scale spatially structured effects.

The Bayesian approach is useful to incorporate expert knowledge and well-known physical behavior about the shape of the functional response of processes to predictor variables into prior models. A large majority of landslides are usually triggered by a specific physical or weather event, such as extreme precipitation in the case of our dataset, and we need spatial random effects to capture the trigger if it has not been observed at high spatial resolution. Moreover, our formulation of the parsimonious space-varying regression is a natural mechanism that allows the model to push the regression coefficient of Slope Steepness toward zero when the trigger is weak. More generally, including such mechanisms into models seems particularly promising for improving interpretability and predictions in cases where a large fraction of the study area experiences only a weak trigger intensity.

Identifiability problems may arise from the use of several latent components, whose spatial resolution and prior specification allows capturing similar types of variability and which are not identifiable in a frequentist framework without prior distributions (e.g., spatial random effects, space-varying regression coefficients and i.i.d. effects, all resolved at the SU level). Typically, the Bayesian paradigm explains spatial variation through the component that is most easily “pushed away” from its prior toward the posterior shape of the point process intensity. Nevertheless, some confounding between such effects is common. For instance, estimated coefficients of fixed effects for covariates whose values are relatively smooth in space often tend to be slightly smaller in absolute value in models with spatially indexed random effects. Despite the complexity of the estimated models, our simulation experiments (see the Supplementary Material, [Opitz et al. \(2022\)](#)) show that there are no strong confounding issues when estimating these models for the studied dataset. If the ultimate goal is spatial prediction for unobserved areas (e.g., prediction of landslide intensity for unobserved SUs adjacent to the observed ones), we must be careful to avoid a transfer of information from spatially dependent components (such as the latent spatial effect in our model) toward spatially independent effects (such as the i.i.d. effect). Our approach relies on penalized complexity priors where we penalize components rather strongly if they lead to very complex models which is an appealing solution to cope with sophisticated latent models that include a moderate number of different random effect components. Moreover, we underscore that the INLA method, combined with modern computing power, provides a convenient toolbox that allows for relatively simple implementation and estimation of very high-dimensional and sophisticated models in the Bayesian framework using latent Gaussian models.

We conclude that model choice is not easy when observed spatial data are discrete, only available through a relatively moderate sample size, not replicated in time, but depend on predictors that may change rapidly even at small spatial scales—in this setting, there usually is not a single “best” model, and data cannot inform us about the true, complex structure of the intensity function of the point process with high certainty. Therefore, careful construction of a moderate number of candidate models using prior expert knowledge is important. We recommend that several model diagnostics related to both goodness-of-fit and predictive

performance be compared and carefully studied in practice. From a pragmatic stance, we emphasize that calculating DIC or WAIC requires only a single estimation of each model. Therefore, they could be used to discard models for which the criterion shows a substantially worse goodness-of-fit than for the group of best-ranking models. For example, our model M1a with pixel-based i.i.d. effect shows a relatively large gap in its DIC and WAIC value with respect to all the other models (see Table 1) and could have been discarded from further analyses based on this result. Predictive comparisons that require repeated estimation of models with some of the data held out, such as k -fold cross-validation, could then have been restricted to a smaller number of models. Since the baseline model considered in this work was already quite complex and highly competitive from a predictive stance (see Lombardo, Opitz and Huser (2018) where a comparison with simpler model structures not studied here is provided), our more sophisticated models could not significantly improve predictive accuracy, as measured through cross-validated predictive scores. However, different model structures provide insightful descriptive analyses, especially when prior physical knowledge is reflected by the model structure, as, for instance, in the parsimonious space-varying regression model M5. Our modeling extensions provided increased interpretability while maintaining the predictive accuracy of simpler reference models. In the literature, the improvement of machine-learning approaches by making their outputs more interpretable and better explainable has been identified as a crucial area of research, and our framework appropriately responds to the three general desiderata formulated for this domain by Murdoch et al. (2019): *predictive accuracy, descriptive accuracy and relevancy of interpretations*, where models should ensure these criteria in the given order of priority.

Acknowledgments. The authors would like to thank two referees, an Associate Editor and the Editor for many valuable comments that helped improve the quality of the manuscript during the review process.

SUPPLEMENTARY MATERIAL

Additional material about data, INLA-based inference and prior sensitivity of results (DOI: [10.1214/21-AOAS1561SUPP](https://doi.org/10.1214/21-AOAS1561SUPP); .pdf). The Supplement contains a detailed description of geomorphological covariates, an introduction to Bayesian estimation with the integrated nested Laplace approximation and results from a simulation study.

REFERENCES

- 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. *Geosci. Model Dev.* **9** 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. *Eng. Geol.* **260** 105237.
- ARNONE, E., FRANCIANE, A., SCARBACI, A., PUGLISI, C. and NOTO, L. (2016). Effect of raster resolution and polygon-conversion algorithm on landslide susceptibility mapping. *Environ. Model. Softw.* **84** 467–481.
- ATKINSON, P. M. and MASSARI, R. (1998). Generalised linear modelling of susceptibility to landsliding in the central Apennines, Italy. *Comput. Geosci.* **24** 373–385.
- AYALEW, L. and YAMAGISHI, H. (2005). The application of GIS-based logistic regression for landslide susceptibility mapping in the Kakuda–Yahiko Mountains, Central Japan. *Geomorphology* **65** 15–31.
- BESAG, J. (1975). Statistical analysis of non-lattice data. *J. Roy. Stat. Soc. (Ser. D)* 179–195.
- BEVEN, K. and KIRKBY, M. J. (1979). A physically based, variable contributing area model of basin hydrology. *Hydrol. Sci. J.* **24** 43–69.
- BOUT, B., LOMBARDO, L., VAN WESTEN, C. J. and JETTEN, V. G. (2018). Integration of two-phase solid fluid equations in a catchment model for flashfloods, debris flows and shallow slope failures. *Environ. Model. Softw.* **105** 1–16.

- CAMA, M., LOMBARDO, L., CONOSCENTI, C., AGNESI, V. and ROTIGLIANO, E. (2015). Predicting storm-triggered debris flow events: Application to the 2009 Ionian Peloritan disaster (Sicily, Italy). *Nat. Hazards Earth Syst. Sci.* **15** 1785–1806.
- CAMA, M., CONOSCENTI, C., LOMBARDO, L. and ROTIGLIANO, E. (2016). Exploring relationships between grid cell size and accuracy for debris-flow susceptibility models: A test in the Giampilieri catchment (Sicily, Italy). *Environmental Earth Sciences* **75** 1–21.
- CARRARA, A., CARDINALI, M., GUZZETTI, F. and REICHENBACH, P. (1995). GIS technology in mapping landslide hazard. In *Geographical Information Systems in Assessing Natural Hazards* 135–175. Springer, Berlin.
- CASTRO CAMILO, D., LOMBARDO, L., MAI, P. M., DOU, J. and HUSER, R. (2017). Handling high predictor dimensionality in slope-unit-based landslide susceptibility models through LASSO-penalized generalized linear model. *Environ. Model. Softw.* **97** 145–156.
- COROMINAS, J., VAN WESTEN, C., FRATTINI, P., CASCINI, L., MALET, J.-P., FOTOPOULOU, S., CATANI, F., VAN DEN EECKHAUT, M., MAVROULI, O. et al. (2014). Recommendations for the quantitative analysis of landslide risk. *Bulletin of Engineering Geology and the Environment* **73** 209–263.
- FAWCETT, T. (2006). An introduction to ROC analysis. *Pattern Recogn. Lett.* **27** 861–874.
- GAMERMAN, D., MOREIRA, A. R. B. and RUE, H. (2003). Space-varying regression models: Specifications and simulation. *Comput. Statist. Data Anal.* **42** 513–533. MR2005406 [https://doi.org/10.1016/S0167-9473\(02\)00211-6](https://doi.org/10.1016/S0167-9473(02)00211-6)
- GELFAND, A. E., KIM, H.-J., SIRMANS, C. F. and BANERJEE, S. (2003). Spatial modeling with spatially varying coefficient processes. *J. Amer. Statist. Assoc.* **98** 387–396. MR1995715 <https://doi.org/10.1198/016214503000170>
- GELMAN, A., HWANG, J. and VEHTARI, A. (2014). Understanding predictive information criteria for Bayesian models. *Stat. Comput.* **24** 997–1016. MR3253850 <https://doi.org/10.1007/s11222-013-9416-2>
- GNEITING, T. and KATZFUSS, M. (2014). Probabilistic forecasting. *Annu. Rev. Stat. Appl.* **1** 125–151.
- GOETZ, J., BRENNING, A., PETSCHKO, H. and LEOPOLD, P. (2015). Evaluating machine learning and statistical prediction techniques for landslide susceptibility modeling. *Comput. Geosci.* **81** 1–11.
- GUZZETTI, F. and REICHENBACH, P. (1994). Towards a definition of topographic divisions for Italy. *Geomorphology* **11** 57–74.
- HEERDEGEN, R. G. and BERAN, M. A. (1982). Quantifying source areas through land surface curvature and shape. *J. Hydrol.* **57** 359–373.
- HUNGR, O., LEROUEIL, S. and PICARELLI, L. (2014). The varnes classification of landslide types, an update. *Landslides* **11** 167–194.
- 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). *Ann. Appl. Stat.* **6** 1499–1530. MR3058673 <https://doi.org/10.1214/11-AOAS530>
- KOH, J., PIMONT, F., DUPUY, J.-L. and OPITZ, T. (2021). Spatiotemporal wildfire modeling through point processes with moderate and extreme marks. ArXiv Preprint. Available at [arXiv:2105.08004](https://arxiv.org/abs/2105.08004).
- KRAINSKI, E. T., 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*. CRC Press/CRC, New York.
- LEININGER, T. J. and GELFAND, A. E. (2017). Bayesian inference and model assessment for spatial point patterns using posterior predictive samples. *Bayesian Anal.* **12** 1–30. MR3597565 <https://doi.org/10.1214/15-BA985>
- 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. *Stoch. Environ. Res. Risk Assess.* **32** 2179–2198.
- LOMBARDO, L., OPITZ, T. and HUSER, R. (2019). Numerical recipes for landslide spatial prediction using R-INLA: A step-by-step tutorial. In *Spatial Modeling in GIS and R for Earth and Environmental Sciences* (H. R. Pourghasemi and C. Gokceoglu, eds.) 55–83. Elsevier, Amsterdam.
- 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., FUBELLI, G., AMATO, G. and BONASERA, M. (2016a). Presence-only approach to assess landslide triggering-thickness susceptibility: A test for the Mili catchment (North-Eastern Sicily, Italy). *Natural Hazards* **84** 565–588.
- LOMBARDO, L., BACHOFER, F., CAMA, M., MÄRKER, M. and ROTIGLIANO, E. (2016b). Exploiting maximum entropy method and ASTER data for assessing debris flow and debris slide susceptibility for the Giampilieri catchment (North-Eastern Sicily, Italy). *Earth Surf. Process. Landf.* **41** 1776–1789.

- LOMBARDO, L., BAKKA, H., TANYAS, H., VAN WESTEN, C., MAI, P. M. and HUSER, R. (2019). Geostatistical modeling to capture seismic-shaking patterns from earthquake-induced landslides. *J. Geophys. Res., Earth Surf.* **124** 1958–1980.
- LOMBARDO, L., OPITZ, T., ARDIZZONE, F., GUZZETTI, F. and HUSER, R. (2020). Space-time landslide predictive modelling. *Earth-Sci. Rev.* 103318.
- MØLLER, J., SYVERSVEEN, A. R. and WAAGEPETERSEN, R. P. (1998). Log Gaussian Cox processes. *Scand. J. Stat.* **25** 451–482. MR1650019 <https://doi.org/10.1111/1467-9469.00115>
- MOORE, I. D., GRAYSON, R. and LADSON, A. (1991). Digital terrain modelling: A review of hydrological, geomorphological, and biological applications. *Hydrol. Process.* **5** 3–30.
- MORAGA, P. (2019). *Geospatial Health Data: Modeling and Visualization with R-INLA and Shiny*. CRC Press/CRC Biostatistics Series, Boca Raton, FL.
- MURDOCH, W. J., SINGH, C., KUMBIER, K., ABBASI-ASL, R. and YU, B. (2019). Definitions, methods, and applications in interpretable machine learning. *Proc. Natl. Acad. Sci. USA* **116** 22071–22080. MR4030584 <https://doi.org/10.1073/pnas.1900654116>
- OPITZ, T. (2017). Latent Gaussian modeling and INLA: A review with focus on space-time applications. *J. French Stat. Soc.* **158** 62–85. MR3720130
- OPITZ, T., HUSER, R., BAKKA, H. and RUE, H. (2018). INLA goes extreme: Bayesian tail regression for the estimation of high spatio-temporal quantiles. *Extremes* **21** 441–462. MR3855716 <https://doi.org/10.1007/s10687-018-0324-x>
- OPITZ, T., BAKKA, H., HUSER, R. and LOMBARDO, L. (2022). Supplement to “High-resolution Bayesian mapping of landslide hazard with unobserved trigger event.” <https://doi.org/10.1214/21-AOAS1561SUPP>
- REICHENBACH, P., ROSSI, M., MALAMUD, B. D., MIHIR, M. and GUZZETTI, F. (2018). A review of statistically-based landslide susceptibility models. *Earth-Sci. Rev.* **180** 60–91.
- ROSSI, M., GUZZETTI, F., REICHENBACH, P., MONDINI, A. C. and PERUCCACCI, S. (2010). Optimal landslide susceptibility zonation based on multiple forecasts. *Geomorphology* **114** 129–142.
- ROUSE JR., J., HAAS, R., SCHELL, J. and DEERING, D. (1974). Monitoring vegetation systems in the Great Plains with ERTS.
- RUE, H. and HELD, L. (2005). *Gaussian Markov Random Fields: Theory and Applications. Monographs on Statistics and Applied Probability* **104**. CRC Press/CRC, Boca Raton, FL. MR2130347 <https://doi.org/10.1201/9780203492024>
- RUE, H., MARTINO, S. and CHOPIN, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **71** 319–392. MR2649602 <https://doi.org/10.1111/j.1467-9868.2008.00700.x>
- RUE, H., RIEBLER, A., SØRBYE, S. H., ILLIAN, J. B., SIMPSON, D. P. and LINDGREN, F. K. (2016). Bayesian computing with INLA: A review. *Annu. Rev. Stat. Appl.* **1**.
- 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. *Statist. Sci.* **32** 1–28. MR3634300 <https://doi.org/10.1214/16-STS576>
- SØRBYE, S. H. and RUE, H. (2014). Scaling intrinsic Gaussian Markov random field priors in spatial modelling. *Spat. Stat.* **8** 39–51. MR3326820 <https://doi.org/10.1016/j.spasta.2013.06.004>
- TENG, M., NATHOO, F. and JOHNSON, T. D. (2017). Bayesian computation for Log-Gaussian Cox processes: A comparative analysis of methods. *J. Stat. Comput. Simul.* **87** 2227–2252. MR3656102 <https://doi.org/10.1080/00949655.2017.1326117>
- TIERNEY, L. and KADANE, J. B. (1986). Accurate approximations for posterior moments and marginal densities. *J. Amer. Statist. Assoc.* **81** 82–86. MR0830567
- 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. *Eng. Geol.* 105942.
- VARNES, D. J. (1958). Landslide types and processes. *Landslides and Engineering Practice* **24** 20–47.
- VRANCKX, M., NEYENS, T. and FAES, C. (2019). Comparison of different software implementations for spatial disease mapping. *Spat. Spatiotemporal. Epidemiol.* **31** 100302. <https://doi.org/10.1016/j.sste.2019.100302>
- WILSON, J. P. and GALLANT, J. C. (2000). Digital terrain analysis. *Terrain Analysis: Principles and Applications* **6** 1–27.
- ZEVENBERGEN, L. W. and THORNE, C. R. (1987). Quantitative analysis of land surface topography. *Earth Surf. Process. Landf.* **12** 47–56.