




Dynamic post-earthquake updating of regional damage estimates using Gaussian Processes

Journal Article

Author(s):

Bodenmann, Lukas ; Reuland, Yves ; Stojadinovic, Bozidar 

Publication date:

2023-06

Permanent link:

<https://doi.org/10.3929/ethz-b-000603080>

Rights / license:

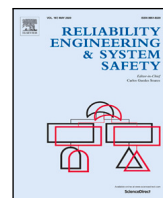
[Creative Commons Attribution 4.0 International](#)

Originally published in:

Reliability Engineering & System Safety 234, <https://doi.org/10.1016/j.ress.2023.109201>

Funding acknowledgement:

821115 - Real-time Earthquake Risk Reduction for Europe (EC)



Dynamic post-earthquake updating of regional damage estimates using Gaussian Processes

Lukas Bodenmann*, Yves Reuland, Božidar Stojadinović

Department of Civil, Environmental and Geomatic Engineering, ETH Zurich, Stefano-Franscini-Platz 5, Zürich, 8093, Switzerland

ARTICLE INFO

Dataset link: <https://doi.org/10.5281/zenodo.7125172>

Keywords:

Post-earthquake damage assessment
Gaussian Process models
Regional earthquake risk models
Uncertainty reduction

ABSTRACT

The widespread earthquake damage to the built environment induces severe short- and long-term societal consequences. Better community resilience may be achieved through well-organized recovery. Decisions to organize the recovery process are taken under intense time pressure using limited, and potentially inaccurate, data on the severity and the spatial distribution of building damage. We propose to use Gaussian Process inference models to fuse the available inspection data with a pre-existing earthquake risk model to dynamically update regional post-earthquake damage estimates and thereby support a well-organized recovery. The proposed method consistently aggregates the gradually incoming building damage inspection data to reduce the uncertainty in ground shaking intensity geographic distribution and to update regional building damage estimates. The performance of the proposed Gaussian Process methodology is demonstrated on one fictitious earthquake scenario and two real earthquake damage datasets. A comparison with purely data-driven methods shows that the proposed method reduces the number of building inspections required to provide reliable and precise damage predictions.

1. Introduction

Damaging earthquakes carry the potential of catastrophic consequences. In addition to direct financial losses, injuries, and fatalities caused by damages to the built environment and short-term disruption in societal functions, earthquakes may give rise to negative long-term consequences, such as increases in home purchase prices and scarcity of low-cost rentals [1]. Housing is an essential infrastructure system, one that provides shelter for its inhabitants, enables higher societal functions and, as shown by the Covid-19 pandemic, provides a safe place to work remotely. The severity and duration of disruptions to such an infrastructure system are not only affected by the seismic resistance of buildings, but also by the efficiency and speed of public and private stakeholder response. Reliable information on the regional severity of damage plays a central role in organizing the emergency response and recovery efforts [2]. However, such information is scarce, incomplete, and imprecise in the early aftermath of an earthquake. Simultaneously, under intense time pressure, stakeholders and agencies need to take crucial decisions that affect immediate disaster assistance as well as the long-term recovery. Currently, visual post-earthquake inspections are conducted in the affected region to enable safe building re-occupancy [3,4]. However, visual inspection of all buildings within the affected region is time- and resource-intensive and may take several

weeks [5,6]. Therefore, decision makers rely on damage estimates deduced from incomplete and uncertain data.

Earthquake risk models may provide initial, quantitative, and spatially exhaustive damage estimates. Such risk models aim to quantify the probability of earthquake-induced consequences on the built environment within a region and, in a broader sense, on its users and inhabitants. While this study focuses on their use for rapid post-earthquake damage estimation [7], risk models are also widely employed in pre-event seismic risk analyses [8,9]. Risk models require information on the exposed buildings, in particular on their seismic vulnerability, together with estimates of the spatial distribution of ground motion intensity. In post-event applications, the latter can be constrained using ground motion recordings [10]. The precision and accuracy of the resulting damage estimates depends, amongst other, on the density of the seismic network, the level of detail of the available exposure information and, finally, how well the seismic vulnerability estimates reflect local conditions of the built inventory. While model-based damage estimates are particularly helpful in the immediate aftermath of an earthquake, when damage observations are not available yet, the inherent multiple sources of uncertainty underline the need for a methodological framework to confirm or correct early damage estimates using empirical evidence, once available.

* Corresponding author.

E-mail addresses: bodenmann@ibk.baug.ethz.ch (L. Bodenmann), reuland@ibk.baug.ethz.ch (Y. Reuland), stojadinovic@ibk.baug.ethz.ch (B. Stojadinović).

Expert-conducted building inspection in the aftermath of an earthquake guarantees a continuous, yet slow, data inflow. Such inspections qualify the induced damage to individual buildings and, as a byproduct, provide important building information that is typically missing in an exposure database (e.g., on the actual lateral load-resisting system). Given this data inflow, machine-learning and statistical inference tools may be leveraged to estimate damage sustained by buildings that have not been inspected yet. Kovačević et al. [11] have applied random forests (RFs) to predict the damage state for individual buildings due to the 2010 Kraljevo (Serbia) earthquake. Stojadinovic et al. [12] have extended the use of RFs and formulated an operational methodology for rapid regional repair cost estimation that requires careful prioritization of inspections. Using inspection data from the 2015 Nepal earthquake, Loos et al. [13] employed geo-statistical techniques to predict the mean damage state aggregated to equally spaced grid cells. However, the proposed method only applies to continuous data, requiring aggregation and averaging of categorical inspection results over grid cells, potentially leading to information loss and constraining the inspection process, because all buildings in a grid cell have to be inspected. Sheibani and Ou [14,15] explored Gaussian Process (GP) models by focusing on continuous quantities that either require sensors within buildings, such as peak floor acceleration, or become available much later in the recovery process (such as repair costs).

We propose to use GP models to fuse the available inspection data with a pre-existing earthquake risk model in order to dynamically update regional post-earthquake damage estimates. The proposed risk model informed GP (RMGP) framework allows processing of observed damage, in the form of multiple ordinal categories (i.e., damage states), to reduce the uncertainty of the geographical distribution of the ground shaking intensity, and to simultaneously update the regional building damage estimates. Thus, instead of learning a new, entirely data-driven, model after the event, individual risk model components are updated. This allows for increasingly constrained predictions of building damage that can be consistently aggregated at different spatial scales. The present work is based on theoretical studies of (approximate) Bayesian updating schemes for infrastructure systems focusing on the failure probability of individual components [16,17] and overall system performance [18–20].

The structure of this paper is: First, we present the mathematical background of GP models in a general context (Section 2), followed by an introduction to earthquake risk models in Section 3. Then, Section 4 describes the risk model inference and updating process using post-earthquake visual damage inspection data of the proposed RMGP framework. Finally, in Section 5, the RMGP framework is applied to the real data from the 1998 Pollino (Italy) and the 2010 Kraljevo (Serbia) earthquakes, as well as to the simulated data for a fictitious event in Zurich (Switzerland).

2. Statistical learning with Gaussian process models

A GP model is a versatile tool to regress and classify data under the Bayesian paradigm [21,22]. GP models assume a directional dependency between a d -dimensional input vector (covariate) \mathbf{x} from some domain \mathcal{X} and the corresponding observable scalar output (response) y . Assuming this dependency is separable into a systematic and a random component, the systematic dependency is given by a latent function $f: \mathcal{X} \rightarrow \mathbb{R}$, such that the likelihood of the output takes the form $p(y|f(\mathbf{x}), \boldsymbol{\vartheta})$, where $\boldsymbol{\vartheta}$ denotes additional parameters of the likelihood. The objective of GP regression and classification is a probabilistic prediction of outputs y_* at the desired target inputs \mathbf{x}_* , via inferring knowledge about the function f from training data and prior beliefs.

Consider a training set of empirical pairwise observations $D = \{(\mathbf{x}_i, y_i) | i = 1, \dots, m\}$, where the vector $\mathbf{y} = [y_1, \dots, y_m]^\top$ and the matrix $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_m]^\top$ collect the outputs and inputs, respectively. Besides the training data, GP inference about f also involves formalizing the prior belief about the latent function. Instead of assuming a fixed

parametric form of f with a finite number of parameters $\boldsymbol{\psi}$ and making inference about $\boldsymbol{\psi}$ only,¹ a GP serves as a prior on the latent function

$$f \sim \mathcal{GP}(m(\mathbf{x}, \boldsymbol{\theta}_m), k(\mathbf{x}, \mathbf{x}', \boldsymbol{\theta}_k)), \quad (1)$$

where $m(\cdot)$ and $k(\cdot)$ denote the mean and positive definite covariance function with parameters $\boldsymbol{\theta}_m$ and $\boldsymbol{\theta}_k$, respectively. While we present mean and covariance functions for regional earthquake risk modeling in Section 3, the reader is referred to [21] for a comprehensive discussion.

A GP is a collection of random variables, any finite subset of which follows a joint multivariate normal distribution. A GP prior on f means that each input vector $\mathbf{x} \in \mathcal{X}$ has an associated random variable $f(\mathbf{x})$, and that the prior joint distribution $p(\mathbf{f}|\mathbf{X}, \boldsymbol{\theta})$ for a collection of function values $\mathbf{f} = [f(\mathbf{x}_1), \dots, f(\mathbf{x}_n)]^\top$ associated with an arbitrary set of n inputs is multivariate normal. The mean vector $\mathbf{m} \in \mathbb{R}^{n \times 1}$ of this distribution has entries $[\mathbf{m}]_i = m(\mathbf{x}_i, \boldsymbol{\theta}_m)$ and the covariance matrix $\mathbf{K} \in \mathbb{R}^{n \times n}$ has entries $[\mathbf{K}]_{ij} = k(\mathbf{x}_i, \mathbf{x}_j, \boldsymbol{\theta}_k)$.

For the inference process, the posterior distribution of the function values \mathbf{f} is computed through Bayes' rule and under the assumption of a factorizing likelihood:

$$p(\mathbf{f}|D, \boldsymbol{\theta}, \boldsymbol{\vartheta}) = \frac{p(\mathbf{y}|\mathbf{f}, \boldsymbol{\vartheta})p(\mathbf{f}|\mathbf{X}, \boldsymbol{\theta})}{p(D|\boldsymbol{\theta}, \boldsymbol{\vartheta})} = \frac{p(\mathbf{f}|\mathbf{X}, \boldsymbol{\theta})}{p(D|\boldsymbol{\theta}, \boldsymbol{\vartheta})} \prod_{i=1}^m p(y_i|f(\mathbf{x}_i), \boldsymbol{\vartheta}), \quad (2)$$

where $p(D|\boldsymbol{\theta}, \boldsymbol{\vartheta}) = \int p(\mathbf{y}|\mathbf{f}, \boldsymbol{\vartheta})p(\mathbf{f}|\boldsymbol{\theta})d\mathbf{f}$ is the marginal likelihood or model evidence. Then, the posterior predictive distribution of function values \mathbf{f}_* at desired target inputs \mathbf{X}_* is derived by marginalizing the function variables over the available training inputs:

$$p(\mathbf{f}_*|D, \boldsymbol{\theta}, \boldsymbol{\vartheta}) = \int p(\mathbf{f}_*|\mathbf{f}, \boldsymbol{\theta})p(\mathbf{f}|D, \boldsymbol{\theta}, \boldsymbol{\vartheta})d\mathbf{f}. \quad (3)$$

The probabilistic predictions of the target outputs are derived by taking the expectation $p(y_*|D, \boldsymbol{\theta}, \boldsymbol{\vartheta}) = \int p(y_*|\mathbf{f}_*, \boldsymbol{\vartheta})p(\mathbf{f}_*|D, \boldsymbol{\theta}, \boldsymbol{\vartheta})d\mathbf{f}_*$. The analytical tractability of posterior distributions depends on the type of the output quantity of interest and assumptions about the likelihood $p(y|f(\mathbf{x}), \boldsymbol{\vartheta})$. The following Section 2.1 serves as an introductory example where we use a Gaussian likelihood. Section 2.2 specifies the non-Gaussian likelihoods required for categorical data.

2.1. Introductory example: Gaussian likelihood

For continuous outputs $y \in \mathbb{R}$ and under the assumption of independent normally distributed noise, we can assume a Gaussian likelihood $p(y|f(\mathbf{x}), \boldsymbol{\vartheta}) = \mathcal{N}(f(\mathbf{x}), \sigma_0)$. The noise variance σ_0^2 is thus an additional parameter of the likelihood $\boldsymbol{\vartheta} = \sigma_0^2$. Bayesian inference of the latent function f in this model is analytically tractable, while the posterior process $f|D$ is again a GP. Thus, the posterior of the function values at the input points of the training set \mathbf{X} , $p(\mathbf{f}|D)$, as well as the posterior predictive distribution for an arbitrary set of target inputs \mathbf{X}_* , $p(\mathbf{f}_*|D)$, follow multivariate normal distributions. By evaluating the covariance matrix between target and training inputs, \mathbf{K}_* , and between the target inputs themselves, \mathbf{K}_{**} , we can compute the mean and covariance of the posterior predictive as:

$$\boldsymbol{\mu}_* = \mathbf{m}_* + \mathbf{K}_*(\mathbf{K} + \sigma_0^2\mathbf{I})^{-1}(\mathbf{y} - \mathbf{m}) \quad (4a)$$

$$\boldsymbol{\Sigma}_{**} = \mathbf{K}_{**} - \mathbf{K}_*(\mathbf{K} + \sigma_0^2\mathbf{I})^{-1}\mathbf{K}_*^\top, \quad (4b)$$

where \mathbf{I} is the identity matrix with the same dimensions as \mathbf{K} .

To illustrate GP regression, we use a one-dimensional toy example, where the prior GP is specified using the popular squared exponential covariance function defined as $k(\mathbf{x}, \mathbf{x}', \boldsymbol{\theta}_k) = \sigma^2 \exp(-\|\mathbf{x} - \mathbf{x}'\|_2^2 / 2\ell^2)$. Here σ^2 is the variance, ℓ the length scale, and $\|\cdot\|_2$ is a Euclidean distance measure. A zero mean function ($m(\mathbf{x}) = 0$) is chosen, limiting the hyper-parameters of the prior GP to $\boldsymbol{\theta} = (\sigma^2, \ell)$.

¹ An example is the linear model $f(\mathbf{x}, \boldsymbol{\psi}) = \mathbf{x}^\top \boldsymbol{\psi}$, where prior uncertainty about f is usually expressed in terms of a prior distribution on $\boldsymbol{\psi}$.

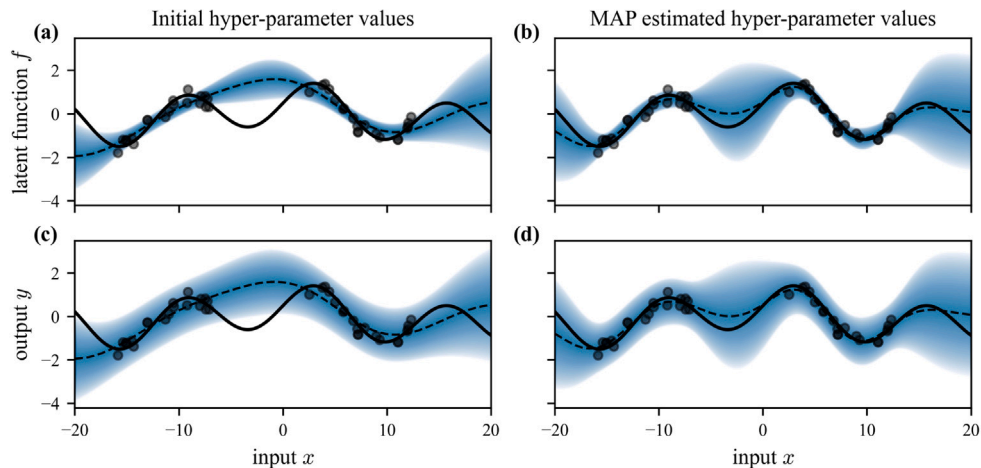


Fig. 1. GP regression with independent normal noise. The input data D , indicated using dots, was generated from the function shown as a solid line with additive normal noise. The top row of plots shows the posterior of the latent functions, where the shaded area shows the density $p(f|D, \theta, \vartheta)$ and the dashed lines indicate the corresponding mean. The bottom row of plots shows the posterior in the output space $p(y|D, \theta, \vartheta)$. In the left column the parameters θ and ϑ are fixed, whereas in the right column the parameters are estimated as their MAP values.

The dataset D consists of $n = 40$ datapoints, at which y is sampled from the function $\sin(0.5x) + 0.5\cos(0.2x)$ polluted with a white noise $\epsilon \sim \mathcal{N}(0, \sqrt{0.1})$. Input positions x are uniformly sampled in two intervals. A first model mimics the situation where the hyper-parameters are fixed to 1, 7 and 0.15 for the GP variance σ^2 , length scale ℓ , and noise variance σ_0^2 , respectively. Fig. 1a shows the posterior predictive distribution over the input space $x_* \in [-20, 20]$. Especially in regions far from available data points, the resulting posterior function gives a poor fit to the generating function, represented by the solid line. Fig. 1c plots the posterior predictive in the output space $p(y_*|D, \theta, \vartheta)$, with the mean function identical to Fig. 1a above, but with the variance being increased by the noise σ_0^2 .

The choice of appropriate hyper-parameters θ and ϑ prior to the inference process is challenging, especially in situations with limited prior domain knowledge. In a full Bayesian setting, hyper-prior distributions $p(\theta)$ and $p(\vartheta)$ are assigned and the joint posterior of the latent function and the hyper-parameters are updated simultaneously. This approach is analytically intractable and thus, requires an approximation via sampling, such as Markov-chain Monte Carlo. We use the maximum a-posteriori (MAP) estimates of the hyper-parameters, for example, the mode of their posterior distributions $p(\theta, \vartheta|D)$, via solving an optimization problem

$$(\hat{\theta}, \hat{\vartheta}) = \underset{\theta, \vartheta}{\operatorname{argmax}} \log p(D|\theta, \vartheta) + \log p(\theta) + \log p(\vartheta), \quad (5)$$

which is analytically tractable for a Gaussian likelihood, and factorizing hyper-prior distributions. In this example we chose log-normal hyper-priors, with the median values identical to the initial values above, and the standard deviations set to 0.8 for ℓ and σ^2 and to 0.2 for σ_0^2 . The posterior function derived using the MAP-estimated hyper-parameters provides a better fit (see Fig. 1b and d), obtained through the reduction of ℓ from 7 to 3. The smaller length scale reduces the overconfidence in regions away from known data points.

GP regression is also widely applied in the geo-statistics field, where it is known as kriging [23] and focuses on two-dimensional input spaces. The method for post-earthquake damage mapping proposed by Loos et al. [13], first uses the non-spatial inputs in least-squares regression to remove a linear trend in the data and secondly uses the geo-coordinates to update GP models with the remaining residuals. The method proposed by Sheibani and Ou [14] also relies on GP regression but applies it to the entire input space and not only on the geo-coordinates. Both methods rely on a Gaussian likelihood which hinders their application to categorical data, such as the damage states assigned during post-earthquake inspections.

2.2. Non-Gaussian likelihoods for categorical data

The next two paragraphs contain commonly used likelihoods for ordered categorical output data (ordinal GP regression) and nominal categorical outputs (multi-class GP classification). Those are of special interest for earthquake risk models, because they describe discrete building damage categories and typological building classes (see Section 3). Finally, a method for approximate Bayesian inference in these non-conjugate cases is presented.

Ordinal GP regression. Ordinal data is defined by scalar outputs y_i being elements of a finite set \mathcal{Y} of $c+1$ ordered categories, which are denoted as integers, $y_i \in \mathcal{Y} = \{0, 1, \dots, c\}$, with preserved ascending ordering information. According to Chu and Ghahramani [24], GP regression can be rewritten for such ordinal data by employing c parameters $\eta_1 < \dots < \eta_c$ to define the class membership probabilities of label y_i conditional on (latent) function value $f_i = f(\mathbf{x}_i)$:

$$p(y_i|f_i, \vartheta) = \begin{cases} 1 - \Phi\left[\frac{f_i - \eta_{y_i+1}}{\beta}\right], & \text{if } y_i = 0 \\ \Phi\left[\frac{f_i - \eta_{y_i}}{\beta}\right], & \text{if } y_i = c \\ \Phi\left[\frac{f_i - \eta_{y_i}}{\beta}\right] - \Phi\left[\frac{f_i - \eta_{y_i+1}}{\beta}\right], & \text{otherwise} \end{cases} \quad (6)$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function. The threshold parameters, $\eta = (\eta_1, \dots, \eta_c)$, partition the real line into contiguous intervals, mapping the continuous function f_i into the discrete variable y_i . The dispersion of these threshold parameters is denoted as β . The parameters of the likelihood function thus become $\vartheta = (\eta, \beta)$.

Multi-class GP classification. Nominal categorical data, similar to ordinal data, is discrete and finite with possible class outputs $y_i \in C = \{0, 1, \dots, c\}$, where c is the number of classes. Unlike ordinal labels, nominal class labels do not provide any ranking information and, instead of having one function f as is the case for ordinal data, we have c independent latent functions, i.e. one for each class. The function values evaluated at data point \mathbf{x}_i are denoted as $\mathbf{g}_i = (g_i^0, \dots, g_i^c)$, where we denote functions as g to avoid confusion with preceding sections. The conditional class membership probabilities of label y_i are calculated using the softmax function [21]:

$$p(y_i|\mathbf{g}_i) = \frac{\exp g_i^{y_i}}{\sum_{k=1}^c \exp g_i^k}. \quad (7)$$

If for class k we write $\mathbf{g}^k = (g_1^k, \dots, g_n^k)$, then $p(\mathbf{g}^k|\mathbf{X})$ follows a multivariate normal distribution. Furthermore, because the c latent functions are independent, so is $p(\mathbf{g}|\mathbf{X})$, where $\mathbf{g} = (\mathbf{g}^1, \dots, \mathbf{g}^c)$. In this case, the likelihood function has no specific parameters $\vartheta = \emptyset$.

Variational Gaussian approximation. Because of the non-Gaussian likelihoods defined in Eq. (6) and Eq. (7), the posterior distribution defined in Eq. (2), is not analytically tractable and we use the variational Gaussian approximation scheme [25] for inference. The posterior is approximated with a multivariate normal distribution $q(\mathbf{f})$ that is close to the true posterior $q(\mathbf{f}) = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \approx p(\mathbf{f}|D, \boldsymbol{\theta}, \boldsymbol{\vartheta})$. Specifically, we search for parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ that minimize the Kullback–Leibler (KL) divergence between the variational distribution $q(\mathbf{f})$ and the true posterior. We can write this KL divergence as

$$\text{KL}(q(\mathbf{f}) \parallel p(\mathbf{f}|D, \boldsymbol{\theta}, \boldsymbol{\vartheta})) = \log p(D|\boldsymbol{\theta}, \boldsymbol{\vartheta}) + \mathcal{F}_{\text{vfe}}(q, \boldsymbol{\theta}, \boldsymbol{\vartheta}), \quad (8)$$

where $\mathcal{F}_{\text{vfe}}(\cdot)$ denotes the variational free energy. Because of the non-Gaussian likelihood, the evidence $p(D|\boldsymbol{\theta}, \boldsymbol{\vartheta})$ is not analytically tractable. However, the evidence being a constant term, minimizing Eq. (8) is equivalent to minimizing the variational free energy. The latter is expressed as

$$\mathcal{F}_{\text{vfe}}(q, \boldsymbol{\theta}, \boldsymbol{\vartheta}) = \text{KL}(q(\mathbf{f}) \parallel p(\mathbf{f}|\boldsymbol{\theta})) - \sum_{i=1}^n \left[\int \log(p(y_i|f_i, \boldsymbol{\vartheta}))q(f_i)df_i \right]. \quad (9)$$

The first term in Eq. (9) is the KL-divergence between two multivariate normal distributions, the variational distribution and the prior of f , and is analytically tractable. The one-dimensional integrals in the second term, also called the expected log likelihood of the data with respect to $q(\mathbf{f})$, can, for example, be approximated using Gauss–Hermite quadrature. In order to keep the variational free energy small, a model should provide a good explanation of the data (via large values for the second term) while not deviating too far from the prior (keeping the first term small). Negative variational free energy is also called the evidence lower bound (ELBO), because it provides a lower bound to the log of evidence. This follows from Eq. (8) and the KL-divergence measure being positive. Therefore, $\mathcal{F}_{\text{vfe}}(\cdot)$ is used to learn the parameters of the variational distribution $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$, while it simultaneously provides approximate MAP estimates for the hyper-parameters $\boldsymbol{\theta}$ and $\boldsymbol{\vartheta}$, using the following objective function:

$$(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}}, \hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\vartheta}}) = \underset{\boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\theta}, \boldsymbol{\vartheta}}{\text{argmax}} - \mathcal{F}_{\text{vfe}}(q, \boldsymbol{\theta}, \boldsymbol{\vartheta}) + \log p(\boldsymbol{\theta}) + \log p(\boldsymbol{\vartheta}). \quad (10)$$

The application of this objective function to derive MAP estimates for the parameters $\boldsymbol{\vartheta}$ of the ordinal likelihood requires a two-fold reparameterization: First, to enable a parallel estimation of the threshold parameters and their dispersion, we perform an up-front scaling $\tilde{\eta} = \eta/\beta$. Second, to fix the order of the threshold parameters, positive parameters $\Delta_y = \tilde{\eta}_y - \tilde{\eta}_{y-1}$ are introduced for categories $y > 0$. Thus, a MAP estimation is performed for parameters $\boldsymbol{\vartheta} = \{\beta, \tilde{\eta}_0, \Delta_1, \dots, \Delta_c\}$.

Because the approximated posterior distribution, $q(\mathbf{f})$, is multivariate normal, the approximate posterior predictive $q(\mathbf{f}_*|\hat{\boldsymbol{\theta}}) = \int p(\mathbf{f}_*|\mathbf{f}, \hat{\boldsymbol{\theta}})q(\mathbf{f})d\mathbf{f}$ is also multivariate normal. Consequently, we can evaluate its parameters as

$$\boldsymbol{\mu}_* = \mathbf{m}_* + \mathbf{A}\boldsymbol{\mu} \quad (11a)$$

$$\boldsymbol{\Sigma}_{**} = \mathbf{A}\boldsymbol{\Sigma}\mathbf{A}^\top + \mathbf{B}, \quad (11b)$$

where $\mathbf{A} = \mathbf{K}_*\mathbf{K}^{-1}$ and $\mathbf{B} = \mathbf{K}_{**} - \mathbf{K}_*\mathbf{K}^{-1}\mathbf{K}_*^\top$. In the following section we link the outlined GP theory to regional earthquake risk models and introduce appropriate mean and covariance functions of the GPs described above.

3. Regional earthquake risk models

Regional risk models, when employed in a post-earthquake context, provide rapid, yet uncertain, damage predictions at regional scale. The geographical distribution of shaking intensity measures (IMs) remains uncertain despite seismic network station measurements. Detailed information about the buildings in the affected region is rarely available. Thus, buildings in the region are clustered into predefined types per their seismic vulnerability. Given the lack of data and knowledge, such

average vulnerability often represents buildings at a national, if not continental, scale [26]. Finally, building type attribution models, or exposure mapping models, correlate building types with socio-economic indicators and building attributes from public databases, such as building height, value and year of construction [27,28], in order to overcome the lack of building-specific information. Such typological attribution models may have limited applicability in the region hit by an earthquake, further adding to the uncertainties arising from the shaking intensity and vulnerability components.

In this study, we assume existence of public databases containing the basic building-specific information. For a building, this information is gathered in a d -dimensional vector \mathbf{x} , whose individual entries are, for example, the geographic coordinates, the construction year, and the number of stories. Building damage is then modeled as a consequence of ground motion intensity at its location, using vulnerability functions expressing the probability of reaching a given damage category as a function of the ground motion IM. We use typological attribution models to map individual buildings to pre-defined building vulnerability types.

This study focuses on the physical damage to residential buildings, because most models of other earthquake consequences, such as recovery time and repair costs, use building damage as the starting point [29,30] and empirical observations of building damage are the main building-specific data source to constrain estimates of other impacts after an earthquake occurs.

3.1. Ground motion intensity measures

Predicting earthquake-induced damage to a spatially distributed set of buildings in this study involves the following assumptions: (i) Conditional on a specific scenario event, the logarithm of the ground motion IM at dispersed building sites follow a multivariate normal distribution; (ii) Conditional on IM, damage to individual buildings is independent; (iii) The capacity of a building is given in terms of an IM threshold at which the building transitions into the next damage state; (iv) The capacities of all buildings in the region are expressed with respect to the same ground motion IM.

Let im_i denote the ground motion IM at location \mathbf{x}_i . The characteristics of the corresponding event are collected in vector \mathbf{e} . To estimate im_i , we employ empirical and ergodic ground motion models (GMMs) of the form

$$\ln im_i = m(\mathbf{x}_i, \mathbf{e}, \boldsymbol{\theta}_m) + \delta_{W,i} + \delta_B, \quad (12)$$

where $m(\cdot)$ denotes the GMM trend function with parameters $\boldsymbol{\theta}_m$ that predicts the median IM conditional on the magnitude, the style of faulting, the source-to-site distance, and the local soil conditions. The terms $\delta_{W,i}$ and δ_B denote the within-event and between-event residuals that are both normally distributed with zero mean and variances σ_W^2 and σ_B^2 , respectively. For a specific event, δ_B is identical for two sites \mathbf{x}_i and \mathbf{x}_j , whereas within-event residuals, $\delta_{W,i}$ and $\delta_{W,j}$ differ. Standard GMMs provide the parameters of Eq. (12) based on mixed-effects regression on ground motion recordings from historic earthquakes.

Past studies suggest that the within-event residuals are spatially dependent, that their joint distribution is multivariate normal [31,32] and that the spatial correlation, ρ , decreases exponentially with the euclidean distance between two sites:

$$\rho(\mathbf{x}_i, \mathbf{x}_j, \ell) = \exp\left(-\frac{\|\mathbf{x}_i^{(1:2)} - \mathbf{x}_j^{(1:2)}\|_2}{\ell}\right), \quad (13)$$

where ℓ denotes the length scale. The superscript (1 : 2) indicates that the correlation function acts on the geographical coordinates (i.e. Easting and Northing) of the input vectors, stored in dimensions one and two.

All random variables in Eq. (12) are Gaussian and the correlation function in Eq. (13) leads to positive definite covariances. Therefore,

the random field of logarithmic IMs $f = \ln im$ follows a GP as defined in Eq. (1). The mean GP function is the same as the trend function of the GMM $m(\mathbf{x}_i, \mathbf{e}, \theta_m)$, while the GP covariance function is $k(\mathbf{x}, \mathbf{x}', \theta_k) = \sigma_W^2 \rho(\mathbf{x}, \mathbf{x}', \ell) + \sigma_B^2$. For a finite set of buildings with inputs \mathbf{X} and event characteristics \mathbf{e} , the joint distribution $p(\mathbf{f}|\mathbf{X}, \mathbf{e}, \theta_f)$ is multivariate normal. The hyper-parameters θ_f of this GP model are the GMM parameters $(\theta_m, \sigma_W^2, \sigma_B^2)$ and the length scale ℓ of the spatial correlation model. The above formulation can be extended to cases involving multiple IMs using separate GPs for each IM. The correlation between the GPs can then be modeled using, for example, a linear model of co-regionalization [33].

3.2. Building damage

We denote the damage state of building i as a discrete random variable Y_i with sequential, collectively exhaustive and mutually exclusive categories $y \in \mathcal{Y} = \{0, 1, \dots, c_y\}$, where $y = 0$ indicates no damage. A building enters damage category $y > 0$, if the ground motion IM at the building location exceeds the corresponding capacity threshold. In compliance with the state of the art, we assume these thresholds to be log-normally distributed with location parameters $\eta_1 < \eta_2 < \dots < \eta_{c_y}$ and a common dispersion parameter β . Based on the logarithm of IM, denoted as f_i , the probability of entering a damage state $y > 0$ is derived as

$$P(Y_i \geq y | f_i, \vartheta) = \Phi \left(\frac{f_i - \eta_y}{\beta} \right), \quad (14)$$

where ϑ collects the parameters of the log-normal distributions. The expression in Eq. (14) is called a vulnerability function. Given the similarity to the likelihood employed in case of ordinal GP regression (Eq. (6)), inference using damage data from a small number of inspected buildings can be performed.

Here we assume that ϑ have been derived for a finite set of vulnerability types $B = \{1, 2, \dots, c_b\}$, such as low-rise unreinforced masonry (URM) or mid-rise reinforced concrete (RC) shear wall buildings. Ideally, the inputs \mathbf{x}_i would contain the type $b_i \in B$ for each building. In reality, however, this is rarely the case, because some building attributes are uncertain, if not unknown, for large subsets of regional building stocks. For example, acquiring input data on the lateral load resisting system often requires detailed and time-consuming field surveys, as described in the following.

3.3. Typological attribution model

Typological attribution, in the present context, describes the process of deriving unknown taxonomic attributes of buildings from known information. The required attribute set depends on the building vulnerability types B ; yet, in our application, the input data \mathbf{x}_i for each building i contains the construction period or year, and the number of floors or an approximate height indicator. The lateral load resisting system (LLRS) is assumed unavailable, as in most real-world scenarios.

We model the randomness pertaining to the LLRS of building i using a discrete random variable A_i with domain $\mathcal{A} = \{1, 2, \dots, c_a\}$. For example, $A_i = 1$ indicates a URM building with flexible floors and $A_i = 2$ indicates an URM building with stiff floors. A typological attribution model provides class membership probabilities for A_i as a function of known input data \mathbf{x}_i . This makes it possible to generate type samples a_i for all buildings first, and then attribute a vulnerability type $b = h(\mathbf{x}, a)$ using a deterministic function $h : \mathcal{X} \cup \mathcal{A} \rightarrow B$, as well as gather the corresponding parameters of the vulnerability function as described above. As an example, the deterministic function $h(\cdot)$ may simply add the height class (i.e., low-rise) of a building, stored in \mathbf{x}_i , to the sampled unknown attribute combination a_i .

In the absence of region-specific data, a typology attribution model based on expert opinion presents the best alternative. Experts may provide insights into the proportions of building types in the region,

for instance, by specifying the evolution of the regional percentage of RC shear wall buildings. In Bodenmann et al. [34], latent functions account for variability and correlations associated to the building attribute combinations to incorporate typological information gathered during post-earthquake inspections. By analogy to the multi-class GP classification scheme (Section 2.2), we employ c_a independent latent functions $\mathbf{g} = [g^{(1)}, \dots, g^{(c_a)}]$ and impose a GP prior on them. The class membership probabilities conditional on these functions are given by Eq. (7). The prior mean function $u(\cdot)$ is calibrated using heuristic proportion estimates. For the covariance function, we combine multiple squared-exponential covariance functions to capture spatially correlated deviations over longer length scales, due to regional differences in construction practices, v_{LS} , and short-scale correlated deviations for classes of nearby buildings that share similar age and number of stories v_{SS} . The resulting covariance function is denoted as

$$v(\cdot) = v_{LS} + v_{SS,1} \cdot v_{SS,2} \cdot v_{SS,3}, \quad (15)$$

where the individual functions of $v_{SS}(\cdot)$ measure similarity with respect to the geo-coordinates, the number of floors and the construction year. The multiplicative structure of the covariance function limits high correlation to buildings, for which all three individual functions are very similar.

3.4. Risk model workflow

The damage to residential buildings \mathbf{X} resulting from a specified earthquake scenario \mathbf{e} is estimated using the following framework: First, we draw r samples from the multivariate normal distributions $p(\mathbf{f}|\mathbf{e}, \mathbf{X})$ (ground motion intensity) and $p(\mathbf{g}|\mathbf{X})$ (typological attribution). Then, for each sample and for each building i , we first sample unknown attributes a_i from the categorical distribution $p(A_i|\mathbf{g}_i)$, with probabilities given by Eq. (7), and subsequently sample a damage state y_i from the categorical distribution $p(Y_i|f_i, a_i)$, with probabilities given by Eq. (6). Finally, for each sample j , buildings with identical state of damage are counted and the values stored in a vector $\mathbf{n}_j = \{n_{y_j}|y \in \mathcal{Y}\}$. Thus, we obtain r samples from $p(\mathbf{N}|\mathbf{e})$, which denotes the joint predictive distribution of how the exposed buildings are partitioned amongst the different damage states. The marginal predictive distributions $p(N_y|\mathbf{e})$ of the entries of random vector \mathbf{N} denote the predicted number of buildings in a certain damage state y .

4. Model updating with post-earthquake data

In the immediate aftermath of an earthquake, it is imperative to use the early-arriving ground motion intensity and building damage data as soon as possible to constrain the uncertainty of regional damage estimates and to organize recovery. The following sections explain the inference steps that use the ground motion intensity data measured by seismic networks, and the building damage and typological data obtained by inspection. Fig. 2 schematically illustrates the proposed risk model informed GP (RMGP) framework. The first damage estimates are obtained using a regional risk model, where estimates of ground motion IMs are constrained by seismic recordings, analogous to the current shake map systems. Then, data from the first inspected buildings is used to further constrain the shake map and, in parallel, update the building vulnerability function parameters as well as the typological attribution model. The Python implementation of the RMGP framework is publicly available [35] and builds upon the software library GPFlow [36].

We use the subscripts S, I and \mathcal{T} to denote index sets related to the seismic stations, the inspected buildings, and the target buildings for which we aim to predict the damage category. In the immediate aftermath of an event, the model predicts the damage to buildings with indices $\mathcal{T} = \{1, 2, \dots, n\}$. Indices of buildings inspected after t time steps are denoted as $I_t \subset \mathcal{T}$, while the indices of the remaining not (yet) inspected buildings are denoted as $\mathcal{T}_t = \mathcal{T} \setminus I_t$.

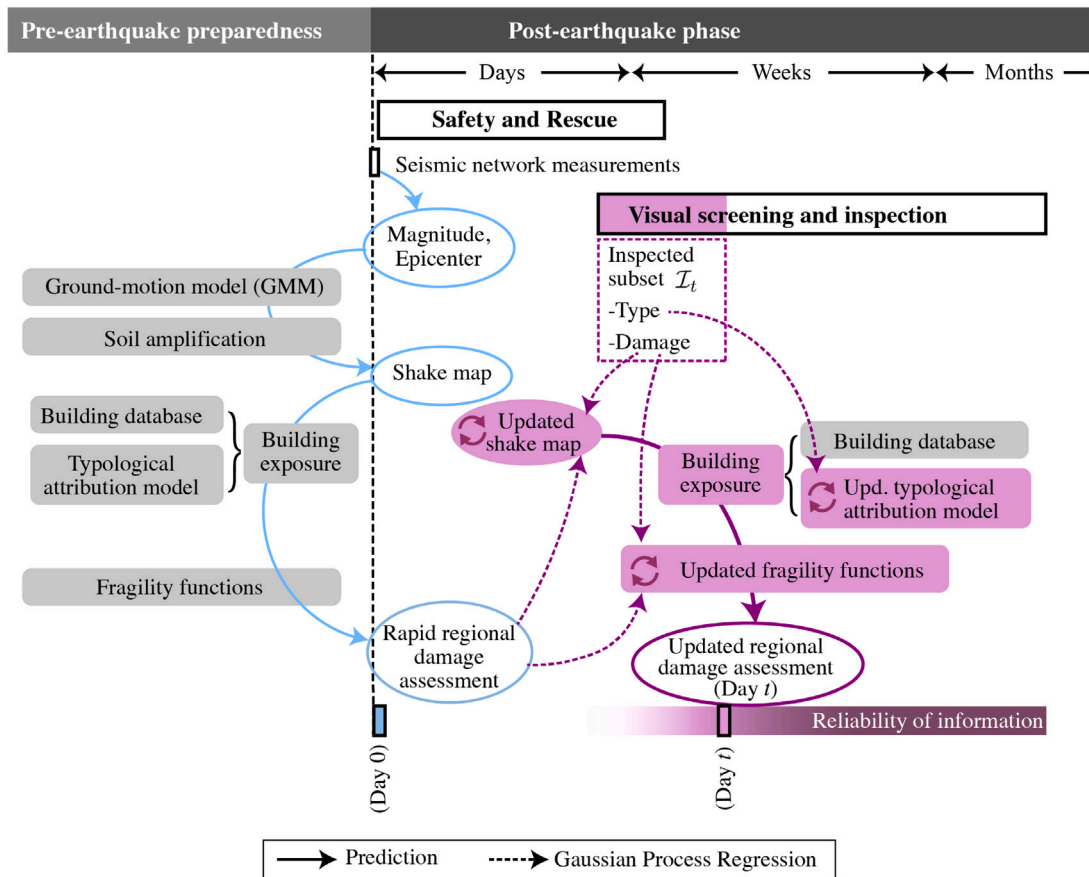


Fig. 2. Schematic overview of the post-earthquake implementation of regional risk models and their GP dynamic updating using inspection data.

4.1. Seismic recordings

Ground motion intensity f is constrained using seismic recordings [37]. We denote the dataset of recorded ground motion IMs from a seismic network as $D_S = \{(x_i, z_i) | i \in S\}$, where z_i are noise-free measurements of the ground motion intensity ($\ln z_i = f_i$). In a first step, we condition the distribution of the (constant) inter-event residual on the seismic recordings by deriving its posterior distribution $p(\delta_B | D_S) = \mathcal{N}(\xi_B, \psi_B)$, where the parameters are calculated as

$$\psi_B^2 = \left(\frac{1}{\sigma_B^2} + \frac{\mathbf{1}^T \mathbf{C}_{SS}^{-1} \mathbf{1}}{\sigma_W^2} \right)^{-1}, \quad (16a)$$

$$\xi_B = \frac{\psi_B^2}{\sigma_W^2} (\mathbf{1}^T \mathbf{C}_{SS}^{-1} (\ln \mathbf{z} - \mathbf{m}_S)), \quad (16b)$$

with \mathbf{C}_{SS} denoting the correlation matrix of the within-event residuals with entries $[\mathbf{C}_{SS}]_{ij} = \rho(\mathbf{x}_{S_i}, \mathbf{x}_{S_j}, \ell)$. Then we estimate the posterior predictive distribution of f evaluated at the target inputs $\mathbf{X}_T = \{\mathbf{x}_i | i \in T\}$, which is a multivariate normal distribution $p(\mathbf{f}_T | D_S) = \mathcal{N}(\mathbf{v}_T, \Psi_{TT})$. The mean vector \mathbf{v}_T and covariance matrix Ψ_{TT} of the posterior predictive distribution are then computed as

$$\mathbf{v}_T = \mathbf{m}_T + \xi_B + \mathbf{C}_{TS} \mathbf{C}_{SS}^{-1} (\mathbf{z} - \mathbf{m}_S - \xi_B), \quad (17a)$$

$$\Psi_{TT} = (\sigma_W^2 + \psi_B^2) (\mathbf{C}_{TT} - \mathbf{C}_{TS} \mathbf{C}_{SS}^{-1} \mathbf{C}_{TS}^T), \quad (17b)$$

where matrices \mathbf{C}_{TS} and \mathbf{C}_{TT} denote the correlation matrices between the target points and seismic stations and between the target points themselves. To generate correlated realizations of damage states for the target buildings \mathbf{X}_T , we follow the workflow described in Section 3.4, where we sample from the posterior $p(\mathbf{f}_T | D_S)$ instead of the prior $p(\mathbf{f}_T)$.

4.2. Inspection data

During inspection of a certain building i experts usually provide an estimate of the inflicted damage y_i via attributing a certain damage category. Here we assume that those are consistent with the damage state \mathcal{Y} employed in the fragility functions of the risk model $y_i \in \mathcal{Y}$. Besides this damage description, experts are also asked to provide some basic building attributes, e.g., the dominating material of the lateral load-resisting system and its type. Again we assume that this information allows attribution of a typological combination $a_i \in \mathcal{A}$ employed in the risk model. We denote the dataset from building inspections conducted up to a given time step t after the event as $D_{I_t} = \{(x_i, y_i, a_i) | i \in I_t\}$, where we omit the subscript t in the following for better readability. This data is used to perform two separate inference steps: First, we perform ordinal GP regression using the function f , where we also account for the seismic recordings D_S . Second, we perform multi-class GP classification using the functions \mathbf{g} related to the unknown typological attributes.

In the first step, we use a_i to allocate a certain fragility class $b_i = h(\mathbf{x}_i, a_i)$ for each of the inspected buildings. Then we use the data y_i and the likelihood for ordinal data (see Eq. (6)) to perform variational inference on f . Specifically, we approximate the true posterior $p(\mathbf{f}_T | D_S, D_{I_t}) \propto p(\mathbf{y} | \mathbf{f}_T) p(\mathbf{f}_T | D_S)$ with a variational multivariate normal distribution $q(\mathbf{f}_T)$. Note that the posterior is conditioned on both the inspection data D_{I_t} and seismic recordings D_S . To evaluate the variational free energy in Eq. (9), we replace the prior $p(\mathbf{f}_T)$ with the posterior predictive $p(\mathbf{f}_T | D_S)$ from Section 4.1. As outlined in Section 2.2, the objective function of this variational inference step (given by Eq. (10)) allows for a simultaneous MAP estimation of the model hyper-parameters. We exploit this capacity in order to update the vulnerability functions. Specifically, we calculate MAP estimates of

Table 1
Characteristics of the earthquake events and the considered regions in three case studies.

	Zurich (simulated)	Pollino (1998)	Kraljevo (2010)
Magnitude M_w	5.8	5.6	5.5
Number of buildings	33594	20528	1959
Area size [km ²]	97.1	931.1	25.4
Number of subregions	22	14	3




Fig. 3. Overview of the three case studies: (a) Simulated M5.8 scenario in the canton of Zurich (Switzerland), (b) Pollino (Italy) earthquake from 1998 and (c) Kraljevo (Serbia) earthquake from 2010. The shaded areas delimit the sub-regions, the star indicates the epicenter location, and the triangles indicate available seismic network stations.

the damage threshold parameters $\hat{\eta}_b$ for all vulnerability types $b \in \mathcal{B}$. Given the limited amount of available building inspections in the early aftermath of an earthquake, we keep the dispersion parameters (β in Eq. (14)) and other hyper-parameters, such as the empirical GMM parameters, fixed at the values specified in the risk model.

In a second step we perform approximate inference on the latent functions \mathbf{g} using the data a_i via standard variational inference as described in Section 2.2. The true posterior $p(\mathbf{g}_I | D_I)$ is approximated by a multivariate normal distribution $q(\mathbf{g}_I)$. We keep the parameters of the mean functions fixed to the initial expert-judgment based estimates and perform MAP estimation for the variances and length scales of the covariance function specified in Eq. (15).

To predict damage for a set of (still) un-inspected target buildings $\mathbf{X}_{\mathcal{T}}$, we follow the procedure outlined in Section 3.4 with two exceptions: First, we sample from the (approximate) posterior predictive distributions $q(\mathbf{f}_{\mathcal{T}})$ and $q(\mathbf{g}_{\mathcal{T}})$. The parameters of these multivariate normal distributions are calculated using Eq. (11). To take into account the seismic recordings in case of function f , we use the mean vector $\mathbf{v}_{\mathcal{T}}$ and covariance matrices $\Psi_{\mathcal{T}\mathcal{T}}$ and $\Psi_{\mathcal{T}I}$ to evaluate matrices \mathbf{A} and \mathbf{B} in Eq. (11). Second, to sample damage categories we employ the updated fragility functions $p(Y|f, a, \vartheta)$, using the inferred MAP-estimates of the threshold parameters.

5. Application to three case-studies

The proposed RMGP framework is applied and evaluated in three case studies, a simulated earthquake scenario in the Swiss Canton of Zurich and two real-world case-studies, the 1998 Pollino (Italy) and the 2010 Kraljevo (Serbia) earthquakes. Specifically, we seek to predict the spatial pattern of earthquake damage as the number of buildings assigned to be in a given damage state within sub-regions (i.e. zip-codes or municipalities). Table 1 summarizes key characteristics of the three examples, while Fig. 3 shows the epicenter locations, the boundaries to the considered sub-regions, and the locations of available seismic network stations. Whereas the earthquake events share similar magnitudes, the case studies strongly differ in terms of the size of the considered region, the number of exposed buildings, and the coverage of the seismic network. The Zurich case represents a densely populated region that is equipped with four seismic network stations. The Pollino

case study covers the largest area but has a relatively low building density. Finally, the Kraljevo case study covers a small region with few buildings and no seismic network stations.

Four states indicate the severity of earthquake damage to residential buildings: none, slight, moderate, and extensive. These four damage states are encoded as integers $\mathcal{Y} = \{0, 1, 2, 3\}$. We mimic a post-event situation with first predictions provided in the immediate aftermath ($t = 0$) using shake maps obtained following the process described in Section 4.1, which are constrained using ground motion recordings D_S . Subsequently, gradually increasing amounts of data from building inspections D_{I_t} enable dynamic updating the risk model components² (see Section 4.2), and thereby lead to improved impact predictions in subsequent time steps $t > 0$. To assess the performance of the proposed RMGP framework in each time step t and for each subregion k , we compare its predictions $N_{k,t}$ of the number of buildings in each damage state to predictions obtained using two entirely data-driven methods (presented in Section 5.2).

5.1. Data and models

The available building-specific input data $\mathbf{X}_{\mathcal{T}}$ consists of the geographical coordinates, the Eurocode 8 soil class [38], the construction year (or period), and the number of stories (either exact or as a range). The event characteristics \mathbf{e} consist of the earthquake (moment-) magnitude, the coordinates of the epicenter and the style of faulting.

The historic inspection data of the Pollino and Kraljevo events do not contain information regarding the inspection sequence. Therefore, the inspection process is simulated in the same manner as for the Zurich case study. Namely, at the beginning of each time step (inspection day), an un-inspected building is randomly assigned to each inspection team. Each inspection team then continues by additionally inspecting four geographically closest (non-inspected) buildings to complete this time step. While producing a spatial clustering of inspection data, this procedure simplifies the time trajectory of the inspection process.

² The inference process at each time step starts from the initial risk model and shake map, and considers data from all buildings that were inspected until this time step.

Therefore, the results are presented as a function of the number of inspected buildings rather than time elapsed since the event.

The supplementary material further describes the three case studies, and details the respective prior risk models, the inspection results, the building databases and, for the Pollino case study, the data pre-processing steps.

5.1.1. Zurich

This case study examines damage to approximately 34000 residential buildings within the Canton of Zurich (Switzerland) inflicted by a $M_w = 5.8$ scenario earthquake (see Fig. 3). Building-specific information $\mathbf{X}_{\mathcal{T}}$ is retrieved from the office for spatial development of the Canton of Zurich (2020). The locations of the seismic network stations $\mathbf{X}_{\mathcal{T}}$ are fictitious and present a dense network. In addition, we take the simplifying assumption of soil class B covering the entire region. The risk model, assumed to be available before the scenario earthquake, builds upon the GMM of Akkar and Bommer [40] and the spatial correlation model of Esposito and Iervolino [41]. We employ vulnerability functions that are defined for the peak ground acceleration (PGA) IM, and specified for 12 vulnerability types.

The data-generating model is deliberately chosen to differ from the above risk model in order to represent real conditions, in which models never fully capture the real processes that govern regional damage distributions. Four cross-correlated random fields of PGA and elastic, 5%-damped spectral acceleration at periods of 0.3, 0.6 and 1.0 s are generated using the GMM of Chiou and Youngs [42]. The spatial cross-correlation of the within-event residuals is accounted for using the model of Markhvida et al. [43] and the cross-correlation of the inter-event residuals follows the relations provided by Baker and Cornell [44]. Finally, the vulnerability functions developed by Martins and Silva [26] are used to simulate damage for each building.

While the available data on coordinates, construction year and number of stories is deemed realistic, the vulnerability functions and soil conditions are not. Therefore, neither the ‘true’ simulated results nor the prior risk model predictions reflect the real earthquake risk conditions in Zurich.

5.1.2. Pollino

On September 9, 1998, the Pollino region in Italy was struck by a $M_w = 5.6$ earthquake that caused widespread building damage. This case study examines approximately 20000 residential buildings spread over 14 subregions (see Fig. 3). The event characteristics and the seismic recordings are taken from the Engineering Strong-Motion database [45]. The main data source for this case study is the publicly available building damage database Da.D.O. [46], containing inspection results for about 13000 residential buildings in the Pollino region. This database provides the geographic coordinates, the construction period (from pre-1919 to post-1981) and the number of stories. We follow Dolce et al. [46], to transform the raw damage descriptions into damage states.

The available data, however, does not cover the entire building stock. The Italian National Institute of Statistics [47] census data from 2001 provides information on the number of residential buildings, their construction period and their number of stories, at municipality level. Based on this information, we augment the original database with the missing buildings, for which spatial coordinates are generated using the European Settlement Map [48]. The 2001 census data is the temporarily closest to the 1998 event, yet, changes that occurred in the three years may induce errors that are deemed small and thus, acceptable. Finally, soil classes are attributed according to Forte et al. [49]. Given the absence of documented damage, the buildings, which are added to augment the database, are assumed to be undamaged. This assumption, despite not reflecting small or undocumented damage, is assumed reasonable.

The employed risk model builds on the GMM of Bindi et al. [50] and the spatial correlation model of Esposito and Iervolino [41]. In

accordance with the Italian risk model [9], buildings are first classified according to their material: URM or RC. Rosti et al. [51,52] empirically derived vulnerability functions for URM and RC buildings, which are used in this case study.

5.1.3. Kraljevo

On November 3, 2010, an earthquake of $M_w = 5.4$ struck the region of Kraljevo (Serbia), claiming two fatalities and leading to reported damage on about 16000 structures. This case study examines damage to approximately 2000 single family residential buildings, for which detailed information are publicly available [12]. We group these 2000 buildings artificially into three subregions (see Fig. 3). The event characteristics are taken from the earthquake catalogue published by the Seismological Survey of Serbia [53]. No seismic recordings from stations sufficiently close to the Kraljevo region are available. The risk model, builds on the GMM of Akkar and Bommer [40] and the spatial correlation model of Esposito and Iervolino [41]. Due to the lack of region-specific vulnerability functions, we attribute functions from Martins and Silva [26], based on the building type descriptions provided in Stojadinovic et al. [12].

5.2. Data-driven methods for comparative analysis

To assess the quality of the predictions obtained using the proposed RMGP framework, two purely data-driven regression methods are used. An ordered linear probit (OLP) model is selected as a simple baseline, while a random forest (RF) model reflects a common approach from literature studies [11,54]. The following building-specific features are used for both approaches: geo-coordinates, epicentral distance, construction year, number of stories, and soil class.

OLP regression is similar to the ordinal GP regression outlined in Section 2.2, with the exception of the latent function f not being a GP but a linear combination of the features $\mathbf{x}^T \boldsymbol{\psi}$. The form of the likelihood is identical to Eq. (6) and the threshold parameters $\boldsymbol{\eta}$, together with the feature weights $\boldsymbol{\psi}$, are estimated by minimizing the negative log-likelihood.

RF classifiers are ensembles of decision trees that are trained using bootstrap samples of the data [55]. For each set of bootstrap samples, a decision tree is grown by recursively partitioning the input space until all terminal nodes contain a pre-defined minimum number of samples. During training, a random subset of the d input variables is selected at each node and the best split is chosen as the split that leads to the largest decrease in Gini-impurity [56]. For prediction, the class membership probabilities are derived via averaging the probabilities predicted by the individual trees. The number of trees in the forest, the minimal required number of samples in a node, and the size of the feature subset to split the nodes are hyperparameters of this model. We fixed the number of trees in the forest to 1000 and perform a grid-search amongst the pre-specified values of the other two hyper-parameters, keeping the combination that returns the highest out-of-bag performance.

5.3. Results

Probabilistic predictions $p(\mathbf{N}_k)$ of the distribution of damage states $y \in \mathcal{Y}$ of the buildings in subregion k are the principal metric to compare the proposed RMGP framework to the purely data-driven models. We introduce two error metrics: (i) the marginal prediction error (MPE) that measures the performance in terms of predicting the number of buildings in any damage state; and (ii) the joint prediction error (JPE) that measures performance jointly over all damage states.

The predictive multivariate distribution $p(\mathbf{N}_k)$ is approximated via r Monte-Carlo samples \mathbf{n}_{kr} , as described in Section 3.4. The vector \mathbf{n}_k denotes for a subregion k the ‘true’ number of buildings in each damage state. We employ the continuous ranked probability score (CRPS) as the error metric to quantify the MPE. The CRPS metric is often used to




Fig. 4. Zurich case study: Evolution of the predicted number of moderately damaged buildings ($y = 2$) in the indicated subregion with increasing amounts of inspection data. Predictions made with the proposed RMGP framework (blue) and an entirely data-driven random forest (RF, orange) model are compared with the ‘true’ value. The numerical values indicate the marginal prediction error (MPE).

Table 2

Marginal and joint prediction error for damage state counts without inspection data ($t = 0$) and using data gathered in three inspection time steps for the subregion indicated in Fig. 4.

			RMGP				RF		
			$t = 0$	$t = 1$	$t = 2$	$t = 3$	$t = 1$	$t = 2$	$t = 3$
Marginal Prediction Error ^a	MPE _{ky}	$y = 0$	118	65	22	26	36	87	27
		$y = 1$	363	15	9	26	58	36	61
		$y = 2$	97	44	15	14	83	115	83
		$y = 3$	14	4	8	4	23	21	16
Joint Prediction Error	JPE _k	365	80	31	41	113	149	107	

^aThe true damage state counts in this subregion are [519, 676, 245, 31] for damage states $y \in \{0, 1, 2, 3\}$.

compare the performance of probabilistic forecasts, for instance from weather forecasting systems [57–59]. Thus, for the amount of buildings in any subregion k being in damage state y , the MPE for predictive distribution $p(N_{ky})$ with respect to the true value \check{n}_{ky} is approximated via r samples as

$$MPE_{ky} \triangleq CRPS(p(N_{ky}); \check{n}_{ky}) \approx \frac{1}{r} \sum_{j=1}^r \left(\|n_{k y j} - \check{n}_{ky}\|_2 - \frac{1}{2r} \sum_{i=1}^r \|n_{k y j} - n_{k y i}\|_2 \right). \quad (18)$$

The first term of the MPE formulation measures the bias - or lack of accuracy - of the predictions, while the second term of the MPE accounts for the dispersion - or lack of precision. For deterministic predictions, the MPE simplifies to the absolute error metric, which renders the interpretation of the MPE more intuitive. Also, using the definition of Eq. (18), the MPE is given in the same units than the predicted quantity, which is, in this case, the number of buildings. Finally, the CRPS can be readily extended to the multivariate case. The energy score (ES) is used to quantify the JPE. The ES generalizes the CRPS towards multivariate predictions. The ES is calculated in a similar way to Eq. (18), by replacing the scalar quantities $n_{k y j}$ and \check{n}_{ky} with vectors \mathbf{n}_{kr} and $\check{\mathbf{n}}_k$, which include all of the damage states. Hence, for the ES, the first term in Eq. (18) measures the Euclidean distance between the sampled damage state partitioning and the true partitioning of the exposed buildings in a four-dimensional space (because we consider four damage states in the case studies). Again, the ES as JPE has the same units as the predicted quantity.

5.3.1. Zurich

The simulated Zurich case study contains the most exhaustive information about the geographical distribution of the ground motion intensity and the ‘true’ damage. Therefore, the local performance of impact estimation is reported in Fig. 4, which contains the predicted

number of moderately damaged buildings in a single subregion after the M5.8 scenario earthquake (see Fig. 3 for the earthquake details). When increasing amounts of buildings are inspected and thus, available to train the GP model, the kernel density of the predictions made with the proposed RMGP framework increases, meaning that the prediction uncertainty decreases. The initial prediction – relying exclusively on the risk model and the seismic-network measurements – suffers from large uncertainties: the amount of buildings with moderate damage ranges from below 100 to over 800 buildings, which corresponds to a range from 7% to 54% of the 1471 buildings in this subregion. Such large uncertainties undermine good decision making and the organization of recovery activities. The predictions stemming from the RF model require a minimum amount of initial building inspections to enable predictions and therefore are represented only for $t \geq 1$. When comparing the proposed RMGP framework with the RF model, the latter returns more precise yet less accurate results. The over-confident predictions with almost no uncertainty produced by the RF model stem from the implicit model assumptions. Conditional on the input values \mathbf{x} of some buildings, the RF model considers damage to these buildings as independent. The RMGP framework, on the other hand, generates building damage conditional on spatially correlated samples of ground motion intensity.

In addition to the probabilistic distribution of the predicted number of moderately damaged buildings, the MPE values, provided in Fig. 4 and derived using Eq. (18), quantify the reduction of the predictive uncertainty and the convergence towards the ‘true’ amount of moderately damaged buildings. Comparing the kernel density distributions with the MPE values highlights the joint contribution of precision and accuracy of the MPE. Providing the MPE for all four damage states (from no damage, $y = 0$, to extensive damage, $y = 3$) of the considered subregion, Table 2 complements Fig. 4. The four time steps coincide with those reported in Fig. 4, where inspection data is available for 0, 0.5, 1, and 1.5% of all 34000 buildings. Finally, Table 2 also provides the JPE,

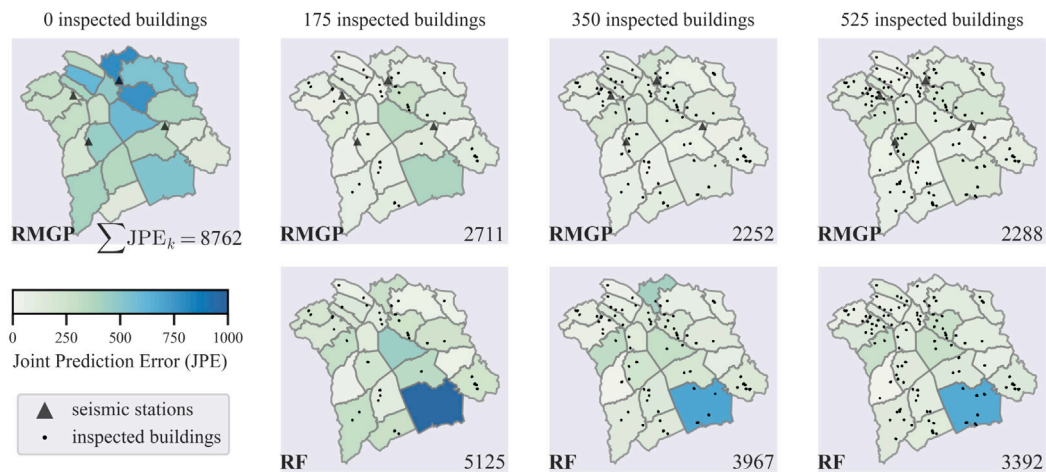


Fig. 5. Zurich case study: JPE for the proposed RMGP framework (top row) and a RF model (bottom row) in all considered subregions at four time steps (columns) with increasing amounts of inspection data. The numerical values indicate the cumulative JPE (over all subregions) achieved by both methods in different time steps.

which measures the error in the joint distribution of building damage states assigned in this subregion.

While Fig. 4 represents the results for a single subregion, impact estimates generally cover the entire region. The damage prediction JPEs for each subregion are reported in Fig. 5. The geographical distribution of prediction errors confirms the rapid convergence of the RMGP framework for all subregions, while the JPE remains high for the RF model in some subregions. Also, the cumulative joint prediction error (CJPE), obtained by summing the JPE over all subregions, shows the convergence of the prediction results towards the ‘true’ building damage.

The CJPE is subsequently used to compare the prediction outcomes of RMGP, RF, and OLP methods in Fig. 6a. The violin outlines provide a visual representation of the distribution of the kernel probability density, i.e., the width of the shaded area corresponds to the proportion of the data at this vertical ordinate. The probability distributions are based on the CJPE of 50 random inspection runs. The dashed and dotted lines within the violins indicate the median and the inter-quartile range.

The RMGP framework produces the most reliable predictions, especially with few available building inspections. The median CJPE of the RMGP framework with 175 inspected buildings is lower than the median CJPE of the RF model after 525 buildings were inspected. Hence, the underlying risk model reduces the need for building inspections in absence of optimized inspection schemes. The influence of the inspection scheme is larger for the purely data-driven methods, RF and OLP.

We use this simulated case study to analyze the performance of the proposed RMGP framework in updating function f , which is the logarithmic ground motion IM, in this case expressed as PGA. The top of Fig. 7 shows the evolution of the inferred median PGA over the entire region when increasing amount of buildings are inspected, while the bottom of the figure illustrates the probability density of the posterior distributions at three locations. The uncertainty at specific locations is reduced, especially when buildings in the vicinity of the location get inspected, as evidenced by the uncertainty reduction after 175 inspections for locations S1 and S2. The variance at location S3 is reduced only after 525 buildings are inspected, some of which are in the location S3 subregion. This comparison is particularly interesting given that the ‘true’ damage for the majority of the inspected buildings originates from IMs other than PGA, namely spectral accelerations at periods of 0.3, 0.6 and 1.0 s.

5.3.2. Pollino

The rapid impact assessment with the risk model for Pollino, shown in Fig. 8 (top left), produces large prediction errors for three subregions.

When inspection outcomes are used to update the risk model, the quality of the estimates increases rapidly in all subregions. The CJPE, representing the total number of mis-predicted buildings, obtained by aggregating subregion data, decreases by over 60%, from 3480 to 1397, by fusing the risk model with inspection data of 600 buildings, corresponding to 3% of the building stock in the region. For this specific inspection sequence, the RF model (shown in the bottom row of Fig. 8) performs better than the initial risk model. In addition, with increasing amounts of data, the prediction performance of the RF model, measured through the CJPE, approaches the one of the RMGP framework.

As outlined in Section 4, the proposed RMGP framework leverages inspection data not only to constrain the distribution of the ground motion IMs, but also to update the vulnerability function parameters. Fig. 9 shows the evolution of the inferred median PGA in the top row, and the evolution of the estimated vulnerability curves for low-rise URM buildings of class A in the bottom row. Compared to the Zurich case study, the Pollino region is almost ten times larger and has a sparse seismic network, which results in larger heterogeneity of the inferred median PGAs. While updating does not change the vulnerability curve for the slight damage state (delimiting $y = 0$ from $y = 1$), the initial vulnerability curve for the extensive damage state shifts towards less conservative predictions after updating using inspection outcomes. Such updated vulnerability curves may be useful in case of aftershocks in the same region, or to develop risk models in other regions with similar building inventories.

The violin outlines of kernel densities of the CJPE of 50 random simulation inspection process outcomes for the Pollino case study are reported in Fig. 6b. Compared with the simulated Zurich earthquake, the initial Pollino risk model is contained within the range of possible outcomes for the RF and the OLP models. Updating the risk model with inspection outcomes in the RMGP framework leads to a median reduction of approximately 50%. In addition, the median and 75-percentile CJPE values for the RMGP framework with 200 inspected buildings are similar to the corresponding RF model values based on 600 building inspections. This underlines the importance of the additional information provided by the risk model in the RMGP framework, especially in the early stages of inspection, when data is scarce.

5.3.3. Kraljevo

The Kraljevo case study covers a much smaller impacted area (see Fig. 3) without seismic network stations. Nevertheless, as seen in Fig. 6c, the predictions originating from the proposed RMGP framework outperform the predictions from the RF and the OLP models, as well as the original risk model predictions. The predictions from RF and OLP models are again very sensitive to changes in the inspection sequence,

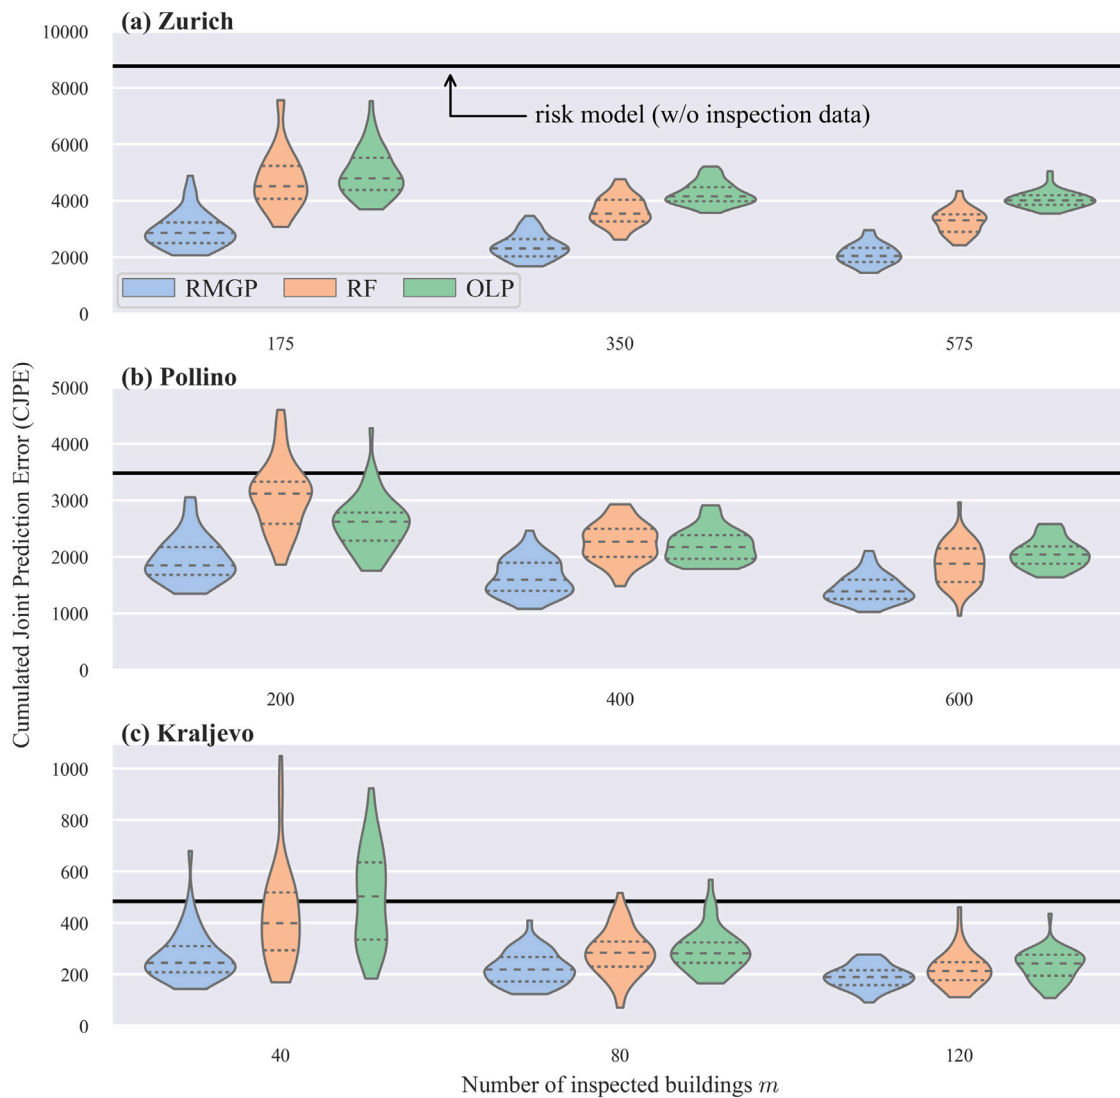


Fig. 6. For the Zurich (a), Pollino (b) and Kraljevo (c) case studies: Violin plots for the cumulative joint prediction error (CJPE) over 50 random inspection processes using the proposed RMGP framework, and the RF and OLP models with increasing amounts of inspection data available for training. The solid black line illustrates the CJPE obtained with the prior risk model and seismic recordings.

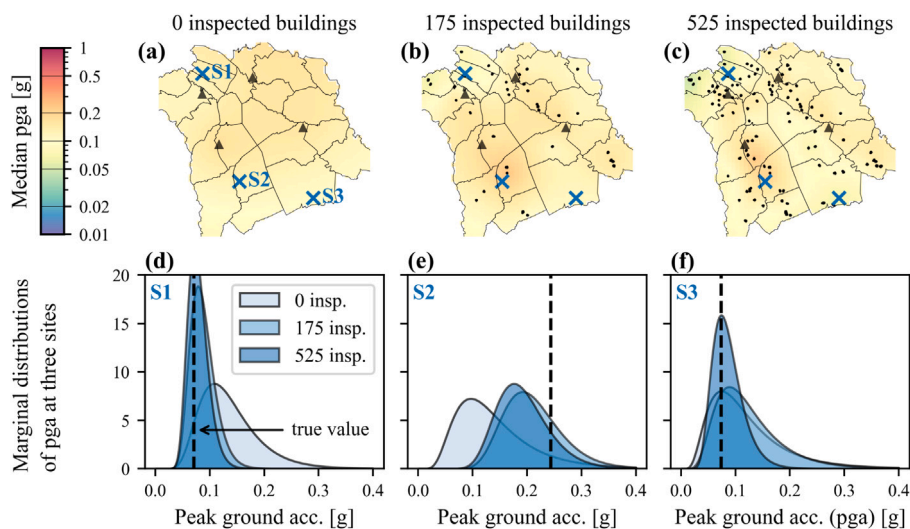


Fig. 7. Zurich case study: Inferred estimates of peak ground acceleration (PGA). The top row (a–c) shows maps of the median PGA inferred from inspection data, while the bottom row (d–f) illustrates the posterior distributions of PGA for three sites S1, S2 and S3.

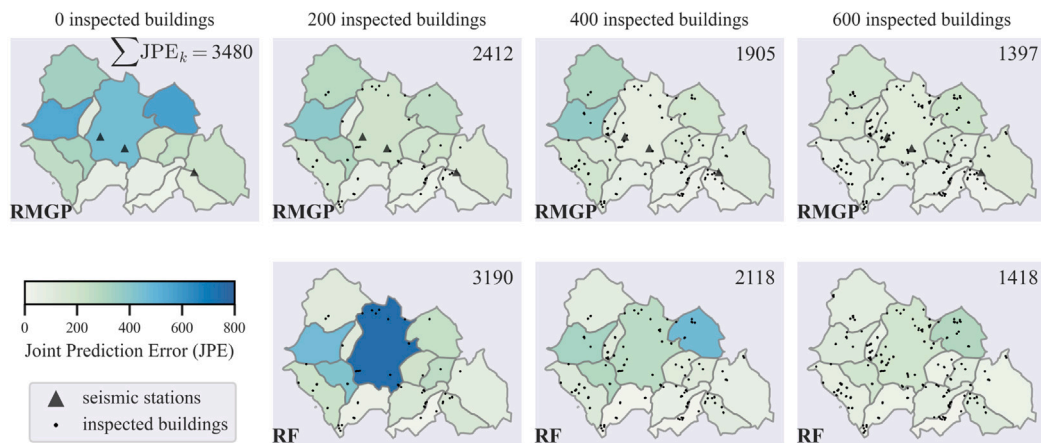


Fig. 8. Pollino case study: JPE for the proposed RMGP framework (top row) and a RF model (bottom row) at four time steps (columns). The numerical values indicate the cumulative (over all subregions) JPE.




Fig. 9. Pollino case study: For increasing amounts of inspection data the top row shows the posterior predictive median PGA, while the bottom row shows the initial and updated vulnerability functions for low-rise URM buildings.

particularly when only 40 building have been inspected. Given that 2000 impacted buildings are considered, data from only 40 building inspections is minuscule, considering the amount of information and correlations to be inferred, especially in absence of an underlying risk model.

5.4. Summary and limitations

Comparing the results of all three case studies, Table 3 summarizes Fig. 6 by reporting the mean, minimum and maximum CJPE values over 50 randomly sampled inspection processes. For a more transparent comparison, the CJPE is normalized by the number of exposed buildings and the best-performing methodology is highlighted in bold for each time step. The potential of the proposed RMGP framework to reduce errors and uncertainty in the geographical distribution of earthquake-induced building damage is confirmed in all three case studies.

The first, simulated, Zurich case study shows that despite a relatively dense seismic network, the initial predictions from the assumed prior risk model may be inaccurate. Updating the components of the risk model using early-arriving building damage inspection data rapidly captures the actual trend in the data and thus, provides reliable regional damage estimates after a fraction of the time that would be required to inspect the entire city. This is also confirmed in the case of Pollino,

where the initial risk-model predictions are more accurate despite the less dense seismic network. While the numbers of inspected buildings at each time step are similar in Zurich and Pollino case studies, we used substantially less inspection data in the case of Kraljevo. The first time step is limited to 40 inspected buildings, which translates to a larger span in CJPEs from different inspection sequences.

Overall, the proposed RMGP framework performs better than purely data-driven models, such as RF. Our focus on the early aftermath, with very few inspected buildings, undermines data-driven methods that, typically, require large training sets. In addition, GP models account for dependencies in the predictions of multiple, spatially distributed outputs. This is particularly important in cases where aggregated damage statistics are of interest. Note that all methods perform inference in less than 40 s on a personal computer and for the maximum number of inspected buildings considered.

It should be noted that the RFs, used for comparison in this paper, are taken from post-earthquake damage prediction literature. Considering the rapid evolution of machine-learning tools, better RF models or models based on other approaches, such as neural networks, may outperform the herein adopted versions of data-driven models. The comparison with the initial rapid damage estimates ($t = 0$) naturally depends on the assumed prior risk models and do not allow for any conclusions on the performance of official national risk models in future earthquake events. Such models might profit from more detailed

Table 3

Summary of the Cumulative Joint Prediction Errors (CJPEs) for all three case studies, using the proposed RMGP framework and the RF and OLP data-drive model, without inspection data ($t = 0$) and using data gathered in three inspection time-steps t after the event. The indicated values are the mean (min., max.) over 50 random inspection processes.

CJPE in % of the number of exposed buildings					
Case study	Model	$t = 0$	$t = 1$	$t = 2$	$t = 3$
Zurich ^a	RMGP	26.1	8.8 (6.1, 14.6)	7.1 (5.0, 10.3)	6.3 (4.3, 8.8)
	RF		14.1 (9.2, 22.5)	10.8 (7.8, 14.2)	9.7 (7.2, 12.9)
	OLP		14.8 (10.6, 22.4)	12.7 (10.6, 15.5)	12.1 (10.6, 15.0)
Pollino ^b	RMGP	17.0	9.5 (6.6, 14.9)	7.9 (5.3, 12.0)	7.0 (5.0, 10.2)
	RF		15.0 (9.1, 22.4)	11.0 (7.2, 14.3)	9.1 (4.7, 14.4)
	OLP		12.7 (8.6, 20.8)	10.7 (8.7, 14.2)	10.0 (8.0, 12.6)
Kraljevo ^c	RMGP	24.7	13.9 (7.3, 34.7)	11.4 (6.3, 20.9)	9.7 (4.7, 14.2)
	RF		21.7 (8.6, 53.5)	14.5 (3.6, 26.4)	11.2 (5.7, 23.5)
	OLP		25.5 (9.4, 47.1)	15.0 (8.4, 29.0)	11.9 (5.5, 22.3)

^aAvailable inspection data for the three time steps: 175 (0.5%), 350 (1.0%), 525 (1.5%).

^bAvailable inspection data for the three time steps: 200 (1.0%), 400 (2.0%), 600 (3.0%).

^cAvailable inspection data for the three time steps: 40 (2.0%), 80 (4.0%), 120 (6.0%).

information on the exposed building stock and can incorporate better expert knowledge.

The dataset of buildings for which inspection data becomes available in the early aftermath is based on random starting points for inspection teams. However, inspections may be performed only on explicit demand from building owners and are thus limited to buildings that have a higher a-priori probability of being damaged than the randomly chosen buildings. To overcome this sampling bias, either random buildings need to be inspected in addition to those with an open inspection demand, or special model training schemes need to be implemented. In the case of the RMGP framework, this could be achieved by conditioning the vulnerability functions on the event that the owner reported building damage.

While this study focuses on conventional post-earthquake data sources (e.g., seismic network stations and building inspections), there are numerous other data sources which provide helpful and important information for post-earthquake response. Satellite and aerial imagery, for example, offer high coverage information that could potentially be combined with the proposed RMGP framework to an ensemble model delivering improved quantitative damage estimates.

6. Conclusion

This paper presents a dynamic updating framework for regional post-earthquake damage estimation, RMGP, that leverages early-arriving observations of ground motion intensity and building damage using Gaussian processes. The framework relies on Gaussian process models to update the ground motion intensity, building type attribution and building type vulnerability components of a regional earthquake risk model that is established prior to the event. Three case studies, focused on regional building damage estimation and updating, illustrate the proposed framework. Our conclusions are:

- RMGP framework provides a powerful tool to combine inspection evidence with prior estimates of traditional risk models, thus enabling precise and accurate predictions of earthquake-inflicted impacts on residential buildings in a fraction of the time required to inspect the entire building stock.
- The parallel updating of ground motion intensity estimates, building vulnerability functions, and building typological attribution models enable a local calibration of prior risk models with broader geographical scopes. The produced byproducts, such as the updated building vulnerability functions, provide useful information for risk modelers beyond rapid damage assessment.
- Using GP models to fuse the early-arriving inspection data with prediction models for probabilistic regional risk analysis requires less data compared to the purely data-driven machine-learning approaches, such as random forest models, to provide these estimates.

- Compared to purely data-driven approaches, predictions obtained using the proposed RMGP framework are more robust with respect to changing inspection sequences and do not require prior optimization or allocation of buildings for priority inspection.

The promising results of the proposed GP dynamic damage updating framework highlight its potential for operational use after future earthquakes. Yet, this requires close collaboration between risk modelers and agencies responsible for planning the inspection process. Future work may involve understanding the effect of different inspection prioritization schemes on the updated predictions, especially in presence of possibly biased inspections, and could further include an active-learning based strategy that optimizes information gain for the RMGP framework. A further challenge lies in effectively combining the dynamically improving damage estimates with damage-to-loss and recovery models to examine the effect of rapid data collection on financial loss predictions and recovery forecasts. This challenge can be extended to include additional interdependent infrastructure systems, such as electric power, potable water or communication. Finally, applying the proposed RMGP framework to more complete data, such as that available at later stages of the inspection process, may offer a robust and automated methodology to derive empirical region-specific building vulnerability curves.

Supplementary material

The code to reproduce the presented case studies is available at <https://doi.org/10.5281/zenodo.7125172> [35]. It includes detailed descriptions of the prior risk models employed in the three case studies, a thorough explanation on the data pre-processing for the Pollino case study, as well as an illustration of the RMGP inference scheme.

CRedit authorship contribution statement

Lukas Bodenmann: Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Visualization, Writing – original draft, Writing – review & editing. **Yves Reuland:** Conceptualization, Visualization, Writing – original draft, , Writing – review & editing. **Božidar Stojadinović:** Resources, Writing – review & editing, Supervision, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The code to reproduce the presented case studies is available at <https://doi.org/10.5281/zenodo.7125172>.

Acknowledgments

The work presented in this paper was financially supported by the Real-time Earthquake Risk Reduction for a Resilient Europe 'RISE' project, financed by the European Union's Horizon 2020 research and innovation program under grant agreement No 821115, as well as the ETH Risk Center project 'DynaRisk', financed under grant agreement ETH-11 18-1.

References

- Potter SH, Becker JS, Johnston DM, Rossiter KP. An overview of the impacts of the 2010–2011 Canterbury earthquakes. *Int J Disaster Risk Reduct* 2015;14:6–14. <http://dx.doi.org/10.1016/j.IJDRR.2015.01.014>.
- Amin S, Goldstein M. Data against natural disasters: Establishing effective systems for relief, recovery, and reconstruction. Washington DC: World Bank; 2008, URL <http://hdl.handle.net/10986/6511>.
- Mcentire DA, Cope J. Damage assessment after the Paso Robles (San Simeon, California) earthquake: Lessons for emergency management. *Natural Hazards Center*; 2004.
- Lallemant D, Soden R, Rubinyi S, Loos S, Barns K, Bhattacharjee G. Post-disaster damage assessments as catalysts for recovery: A look at assessments conducted in the wake of the 2015 Gorkha, Nepal, Earthquake. *Earthq Spectr* 2017;33(Special issue 1):S435–51. <http://dx.doi.org/10.1193/120316EQS222M>.
- Marquis F, Kim JJ, Elwood KJ, Chang SE. Understanding post-earthquake decisions on multi-storey concrete buildings in Christchurch, New Zealand. *Bull Earthq Eng* 2017;15(2):731–58. <http://dx.doi.org/10.1007/S10518-015-9772-8/FIGURES/5>.
- Kusunoki K. Damage assessment in Japan and potential use of new technologies in damage assessment. In: *Advances in assessment and modeling of earthquake loss*. 2021, p. 27–46. <http://dx.doi.org/10.1007/97830306881342>.
- Guérin-Marthe S, Gehl P, Negulescu C, Auclair S, Fayjaloun R. Rapid earthquake response: The state-of-the-art and recommendations with a focus on European systems. *Int J Disaster Risk Reduct* 2021;52:101958. <http://dx.doi.org/10.1016/J.IJDRR.2020.101958>.
- Silva V, Amo-Oduro D, Calderon A, Costa C, Dabbeek J, Despotaki V, et al. Development of a global seismic risk model. *Earthq Spectr* 2020;36(1_suppl):372–94. <http://dx.doi.org/10.1177/8755293019899953>.
- Dolce M, Protà A, Borzi B, da Porto F, Lagomarsino S, Magenes G, et al. Seismic risk assessment of residential buildings in Italy. *Bull Earthq Eng* 2021;19(8):2999–3032. <http://dx.doi.org/10.1007/s10518-020-01009-5>.
- Wald DJ, Quitoriano V, Heaton TH, Kanamori H, Scrivner CW, Worden CB. TriNet “ShakeMaps”: Rapid generation of peak ground motion and intensity maps for earthquakes in Southern California. *Earthq Spectr* 1999;15(3):537–55. <http://dx.doi.org/10.1193/1.1586057>, URL <http://journals.sagepub.com/doi/10.1193/1.1586057>.
- Kovačević M, Stojadinović Z, Marinković D, Stojadinović B. Sampling and machine learning methods for a rapid earthquake loss assessment system. In: *11th U.S. national conference on earthquake engineering*. Los Angeles, CA: Earthquake Engineering Research Institute; 2018.
- Stojadinovic Z, Kovacevic M, Marinkovic D, Stojadinovic B. Rapid earthquake loss assessment based on machine learning and representative sampling. *Earthq Spectr* 2021. <http://dx.doi.org/10.1177/87552930211042393>.
- Loos S, Lallemant D, Baker J, McCaughey J, Yun SH, Budhathoki N, et al. G-DIF: A geospatial data integration framework to rapidly estimate post-earthquake damage. *Earthq Spectr* 2020;36(4):1695–718. <http://dx.doi.org/10.1177/875529302020926190>.
- Sheibani M, Ou G. The development of Gaussian process regression for effective regional post-earthquake building damage inference. *Comput-Aided Civ Infrastruct Eng* 2021;36(3):264–88. <http://dx.doi.org/10.1111/MICE.12630>.
- Sheibani M, Ou G. Adaptive local kernels formulation of mutual information with application to active post-seismic building damage inference. *Reliab Eng Syst Saf* 2021;215. <http://dx.doi.org/10.1016/j.res.2021.107915>.
- Pozzi M, Wang Q. Gaussian process regression and classification for probabilistic damage assessment of spatially distributed systems. *KSCSE J Civ Eng* 2018;22(3):1016–26. <http://dx.doi.org/10.1007/s12205-018-0014-x>.
- Straub D, Der Kiureghian A. Improved seismic fragility modeling from empirical data. *Struct Saf* 2008;30(4):320–36. <http://dx.doi.org/10.1016/j.strusafe.2007.05.004>.
- Bensi M, Der Kiureghian A, Straub D. Bayesian network modeling of correlated random variables drawn from a Gaussian random field. *Struct Saf* 2011;33(6):317–32. <http://dx.doi.org/10.1016/j.strusafe.2011.05.001>.
- Gehl P, Cavalieri F, Franchin P. Approximate Bayesian network formulation for the rapid loss assessment of real-world infrastructure systems. *Reliab Eng Syst Saf* 2018;177:80–93. <http://dx.doi.org/10.1016/j.res.2018.04.022>.
- DeJesus Segarra J, Bensi M, Modarres M. A Bayesian network approach for modeling dependent seismic failures in a nuclear power plant probabilistic risk assessment. *Reliab Eng Syst Saf* 2021;213. <http://dx.doi.org/10.1016/j.res.2021.107678>.
- Rasmussen C, Williams K. *Gaussian processes for machine learning*. Cambridge, MA: The MIT Press; 2006.
- Kuss M. *Gaussian process models for robust regression, classification, and reinforcement learning* [Ph.D. thesis], Darmstadt, Germany: Technische Universität Darmstadt; 2006.
- Matheron G. The intrinsic random functions and their applications. *Adv Appl Probab* 1973;5(3):439–68. <http://dx.doi.org/10.2307/1425829>.
- Chu W, Ghahramani Z. Gaussian processes for ordinal regression. *J Mach Learn Res* 2005;6:1019–41, URL <http://www.jmlr.org/papers/volume6/chu05a/chu05a.pdf>.
- Opper M, Archambeau C. The variational Gaussian approximation revisited. *Neural Comput* 2009;21(3):786–92. <http://dx.doi.org/10.1162/neco.2008.08-07-592>.
- Martins L, Silva V. Development of a fragility and vulnerability model for global seismic risk analyses. *Bull Earthq Eng* 2020. <http://dx.doi.org/10.1007/s10518-020-00885-1>.
- Diana L, Thiriot J, Reuland Y, Lestuzzi P. Application of association rules to determine building typological classes for seismic damage predictions at regional scale: The case study of Basel. *Front Built Environ* 2019;5. <http://dx.doi.org/10.3389/fbuil.2019.00051>.
- Crowley H, Despotaki V, Rodrigues D, Silva V, Toma-Danila D, Riga E, et al. Exposure model for European seismic risk assessment. *Earthq Spectr* 2020;36(1_suppl):252–73. <http://dx.doi.org/10.1177/8755293020919429>.
- Burton HV, Miles SB, Kang H. Integrating performance-based engineering and Urban simulation to model post-earthquake housing recovery. *Earthq Spectr* 2019;34(4):1763–85. <http://dx.doi.org/10.1193/041017EQS067M>.
- Blagojević N, Didier M, Stojadinović B. Quantifying component importance for disaster resilience of communities with interdependent civil infrastructure systems. *Reliab Eng Syst Saf* 2022;228. <http://dx.doi.org/10.1016/j.res.2022.108747>.
- Goda K, Hong HP. Spatial correlation of peak ground motions and response spectra. *Bull Seismol Soc Am* 2008;98(1):354–65. <http://dx.doi.org/10.1785/0120070078>.
- Jayaram N, Baker JW. Statistical tests of the joint distribution of spectral acceleration values. *Bull Seismol Soc Am* 2008;98(5):2231–43. <http://dx.doi.org/10.1785/0120070208>.
- Loth C, Baker JW. A spatial cross-correlation model of spectral accelerations at multiple periods. *Earthq Eng Struct Dyn* 2013;42(3):397–417. <http://dx.doi.org/10.1002/eqe.2212>.
- Bodenmann L, Reuland Y, Stojadinovic B. Dynamic updating of building loss predictions using regional risk models and conventional post-earthquake data sources. In: *Proceedings of the 31st European safety and reliability conference*. 2021, p. 1411–8. [http://dx.doi.org/10.3850/978-981-18-2016-8\(\)559-cd](http://dx.doi.org/10.3850/978-981-18-2016-8()559-cd).
- Bodenmann L. Dynamic updating of post-earthquake loss estimates using Risk Model informed Gaussian Processes (RMGP) (v1.1.0). 2023, <http://dx.doi.org/10.5281/zenodo.7125172>, Zenodo [software].
- Matthews A, van der Wilk M, Nickson T, Fujii K, Boukouvalas A, León-Villagrà P, et al. GPflow: A Gaussian process library using TensorFlow. *J Mach Learn Res* 2017;18:1–6, URL <http://jmlr.org/papers/v18/16-537.html>.
- Worden C, Thompson EM, Baker JW, Bradley BA, Luco N, Wald DJ. Spatial and spectral interpolation of ground-motion intensity measure observations. *Bull Seismol Soc Am* 2018;108(2):866–75. <http://dx.doi.org/10.1785/0120170201>.
- Eurocode 8. *Design of structures for earthquake resistance - Part 1: General rules, seismic actions and rules for buildings*. Brussels, Belgium: European Committee for Standardization (CEN); 2004.
- Amt für Raumentwicklung. Geografisches Informationssystem des Kantons Zürich (GIS-ZH). 2020, GIS-Browser, URL <http://maps.zh.ch>.
- Akkar S, Bommer JJ. Empirical equations for the prediction of PGA, PGV, and spectral accelerations in Europe, the Mediterranean Region, and the Middle East. *Seismol Res Lett* 2010;81(2):195–206. <http://dx.doi.org/10.1785/gssrl.81.2.195>.
- Esposito S, Iervolino I. PGA and PGV spatial correlation models based on European multievent datasets. *Bull Seismol Soc Am* 2011;101(5):2532–41. <http://dx.doi.org/10.1785/0120110117>.
- Chiou BS-J, Youngs RR. Update of the chiou and Youngs NGA model for the average horizontal component of peak ground motion and response spectra. *Earthq Spectr* 2014;30(3):1117–53. <http://dx.doi.org/10.1193/072813EQS219M>.
- Markhvida M, Ceferino L, Baker JW. Modeling spatially correlated spectral accelerations at multiple periods using principal component analysis and geostatistics. *Earthq Eng Struct Dyn* 2018;47(5):1107–23. <http://dx.doi.org/10.1002/eqe.3007>.
- Baker JW, Cornell A. Correlation of response spectral values for multicomponent ground motions. *Bull Seismol Soc Am* 2006;96(1):215–27. <http://dx.doi.org/10.1785/0120050060>.

- [45] Luzi L, Lanzano G, Felicetta C, D'Amico M, Russo E, Sgobba S, et al. Engineering strong motion database (ESM) (Version 2.0). Istituto Nazionale di Geofisica e Vulcanologia (INGV); 2020, <http://dx.doi.org/10.13127/ESM.2>.
- [46] Dolce M, Speranza E, Giordano F, Borzi B, Bocchi F, Conte C, et al. Observed damage database of past Italian earthquakes: The Da.D.O. WebGIS. *Bollettino Geofisica Teorica Appl* 2019;60(2):141–64. <http://dx.doi.org/10.4430/bgta0254>.
- [47] ISTAT. Censimento generale della popolazione e delle abitazioni. Italian National Institute of Statistics; 2001, URL <http://dawinci.istat.it/jsp/MD/>.
- [48] Corbane C, Sabo F. European settlement map from copernicus very high resolution data for reference year 2015, public release 2019. European Commission, Joint Research Centre (JRC); 2019, <http://dx.doi.org/10.2905/8BD2B792-CC33-4C11-AFD1-B8DD60B44F3B>.
- [49] Forte G, Chioccarelli E, De Falco M, Cito P, Santo A, Iervolino I. Seismic soil classification of Italy based on surface geology and shear-wave velocity measurements. *Soil Dyn Earthq Eng* 2019;122:79–93. <http://dx.doi.org/10.1016/j.soildyn.2019.04.002>.
- [50] Bindi D, Pacor F, Luzi L, Puglia R, Massa M, Ameri G, et al. Ground motion prediction equations derived from the Italian strong motion database. *Bull Earthq Eng* 2011;9(6):1899–920. <http://dx.doi.org/10.1007/s10518-011-9313-z>.
- [51] Rosti A, Rota M, Penna A. Empirical fragility curves for Italian URM buildings. *Bull Earthq Eng* 2021;19(8):3057–76. <http://dx.doi.org/10.1007/s10518-020-00845-9>.
- [52] Rosti A, Del Gaudio C, Rota M, Ricci P, Di Ludovico M, Penna A, et al. Empirical fragility curves for Italian residential RC buildings. *Bull Earthq Eng* 2021;19(8):3165–83. <http://dx.doi.org/10.1007/s10518-020-00971-4>.
- [53] Seismological Survey of Serbia. National earthquake catalogue. Technical report, 2017, URL <http://www.seismo.gov.rs/Seizmicnost/Katalog-zemljotresa.pdf>.
- [54] Roeslin S, Ma Q, Juárez-García H, Gómez-Bernal A, Wicker J, Wotherspoon L. A machine learning damage prediction model for the 2017 Puebla-Morelos, Mexico, earthquake. *Earthq Spectr* 2020;36(2_suppl):314–39. <http://dx.doi.org/10.1177/8755293020936714>.
- [55] Breiman L. Random forests. *Mach Learn* 2001;45:5–32. <http://dx.doi.org/10.1023/A:1010933404324>.
- [56] Breiman L, Friedman JH, Olshen RA, Stone CJ. Classification and regression trees. 1984, <http://dx.doi.org/10.1201/9781315139470>.
- [57] Hersbach H. Decomposition of the continuous ranked probability score for ensemble prediction systems. *Weather Forecast* 2000;15(5):559–70. [http://dx.doi.org/10.1175/1520-0434\(2000\)015<0559:DOTCRP>2.0.CO;2](http://dx.doi.org/10.1175/1520-0434(2000)015<0559:DOTCRP>2.0.CO;2).
- [58] Gneiting T, Raftery AE. Strictly proper scoring rules, prediction, and estimation. *J Amer Statist Assoc* 2007;102(477):359–78. <http://dx.doi.org/10.1198/016214506000001437>.
- [59] Gneiting T, Katzfuss M. Probabilistic forecasting. *Annu Rev Stat Appl* 2014;1(1):125–51. <http://dx.doi.org/10.1146/annurev-statistics-062713-085831>.