KSCE Journal of Civil Engineering (2018) 22(3):1016-1026

DOI 10.1007/s12205-018-0014-x

Probabilistic Learning for Decision-making on Civil Infrastructures

pISSN 1226-7988, eISSN 1976-3808
www.springer.com/12205

Gaussian Process Regression and Classification for Probabilistic Damage Assessment of Spatially Distributed Systems

Matteo Pozzi* and Qiaochu Wang**

Received November 23, 2017/Revised December 20, 2017/Accepted December 25, 2017/Published Online

Abstract

We illustrate how methods for non-parametric regression and classification based on Gaussian Processes can be adapted for inferring the condition state of infrastructure components under spatially distributed stressors. When the stressor is modeled by a random field, observations collected in one location can reduce the uncertainty about the stressor intensity also in other locations. Exact inference is possible when the field is Gaussian and observations are perfect or affected by Gaussian noise. However, often the available observations are binary, as those related to the failure or survival of components, and indicate whether the local field is above or below a threshold whose value may also be uncertain. While no efficient scheme for exact inference is available in that setting, we can perform efficient approximate inference when the field is Gaussian and so is the uncertainty on the threshold value. The mathematical formulation for this problem is analogous to that of classification in machine learning, that can be based on latent Gaussian processes. We show how to formulate the problem and how to adapt deterministic methods, as Laplace's method and Expectation Propagation, and methods based on random number generation, as Monte Carlo uniform sampling and importance sampling, to perform approximate inference. Our illustrative application is the condition assessment of assets exposed to a seismic event. Under specific assumptions, the seismic demand can be modeled as a Gaussian random field, and measures about the demand and about the survival and failure of assets can be processed globally, to update the risk assessment. Specifically, we evaluate methods for approximate inference and discuss their merits.

Keywords: *binary components, gaussian process classification, probabilistic inference, spatially distributed systems*

1. Introduction

Observations collected in-field can have a significant impact in the condition assessment of infrastructure components. When physical values at different places are statistically interdependent, information collected at one location is relevant also for others. Bayesian analysis provides a consistent framework for this inference process, propagating information across space. In this framework, we define the joint probability of the set of variables related to different locations, and derive conditional probabilities depending on the observed values. Computationally, inference among a high number of variables is feasible only under specific assumptions. One of these settings requires modeling a spatially distributed stressor with a Gaussian random field, which is completely defined by its mean and covariance functions. Exact inference for these fields is possible for perfect observations and for observations affected by Gaussian noise. This is the traditional setting for Kriging (Krige, 1951) and Gaussian Process nonparametric regression (Rasmussen and Williams, 2006), which has been recently adapted to the monitoring of infrastructure systems under seismic risk (Malings and Pozzi, 2016a). However, even when the stressor is modeled by a Gaussian random field, we cannot effectively perform exact inference when observations are related to non-Gaussian likelihood functions. A case of interested is the following: suppose we can observe the value of binary variables, whose state depends on the stressor being above or below a threshold. Possibly, the threshold value is also unknown, and we model this uncertainty by using a Gaussian variable. In infrastructure condition assessment, this is the case of observing the condition states (e.g. failed or intact) of a set of components in certain locations, with the goal of predicting the condition of uninspected components in other locations. If these conditions depend on a distributed field, we can update the field model by processing those binary observations, and base the prediction on the updated model. For example, observing that some components failed under seismic excitation in a certain area should increase the posterior probability of failure of similar components in a location nearby, as the seismic intensity is likely to be high in the observed area and, being the field smooth, also be high in the location nearby. However, the likelihood function for binary observations is not Gaussian, nor it is the exact posterior field, and no simple inference approach for conjugate

*Associate Professor, CEE Dept., Carnegie Mellon University, Pittsburgh, PA, USA (Corresponding Author, E-mail: mpozzi@cmu.edu)
**M.S. Student, CEE Dept., Carnegie Mellon University, Pittsburgh, PA, USA (E-mail: wangqiaochu1994@gmail.com)

distributions can be used. One approach to approximate inference is to use Monte Carlo methods to generate samples, possibly weighted, for representing the posterior distribution. Uniform Sampling, Importance Sampling or Monte Carlo Markov Chain (MacKay, 2003) can be used for this purpose. As an alternative, we can adopt a deterministic method to identify the Gaussian distribution that best approximates the posterior one. The problem is analogous to that of binary classification in machine learning (Rasmussen and Williams, 2006), where a binary classifier is trained by processing a labeled dataset (e.g. for distinguishing between ham and spam emails). Latent Gaussian Processes provide a flexible non-parametric approach for this classification problem, and different methods for deterministic approximate inference have been proposed. In this paper we illustrate how, under certain conditions, the probabilistic condition assessment of spatially distributed systems can be formulated in analogy with the regression/classification problem in machine learning, adapting and generalizing the formulation of Rasmussen and Williams (2006). We investigate the performance of approximate methods as Laplace's Method and Expectation Propagation, by comparing them with Uniform and Importance Sampling Monte Carlo.

Bayesian inference for spatially distributed systems under seismic risk, applied to the processing of structural condition data, has been presented by Yue *et al.* (2012) and Pozzi and der Kiureghian (2013), both using Gaussian variables, and by Bensi *et al.* (2014) and Cavalieri *et al.* (2016), discretizing the field to perform inference using discrete variables. This paper focuses on methods allowing the process of large datasets, to predict the condition of large sets of uninspected components.

2. Problem Statement

2.1 The Engineering Problem

Consider an infrastructure system, or a set of assets, spatially distributed in a region, exposed to a common stressor. We collect two types of measurements: N_y local observations of stressor intensity and N_s observations of the condition of components hit by the stressors. Intensity measures can be affected by noise, while the condition state depends on the demand posed by the stressor and on the component's capacity. The task is to infer the spatial intensity map of the stressor and predict the damage of other N_* uninspected components. When the intensity is assumed to be smooth in space, the observation of some components' state should affect the risk for components with similar capacities in the proximity of the observed ones: the engineering problem is how to consistently update the risk, by processing observations of intensity and condition state.

2.2 Mathematical Formulation

Let us model the intensity f of a spatially distributed stressor affecting a region with a Gaussian random field:

$$f(\mathbf{x}) \sim \mathcal{GP}\left[\mu_f(\mathbf{x}), k_f(\mathbf{x}, \mathbf{x}')\right] \tag{1}$$

where $\mathbf{x} \in \mathcal{D} \subseteq \mathbb{R}^n$ identifies a point in the n dimensional domain $\mathcal{D}$, $\mathcal{GP}$ indicates a Gaussian Process, μ_f is the mean function and k_f the covariance function, defined for each pair of points, $\mathbf{x}$ and $\mathbf{x}'$, in $\mathcal{D}$. For a 2-D geographical region, n is two. We have collected N_y noisy observations of the intensity. Observation i has been collected in location $\mathbf{x}_i$, where demand is $f_i = f(\mathbf{x}_i)$, and measure is $y_i = f_i + \varepsilon_i$. For the sake of simplicity, we can assume a homoscedastic noise $\varepsilon_i \sim \mathcal{N}(0, \sigma_\varepsilon^2)$, independent respect to all noise values affecting other measures, and to the intensity field, $\mathcal{N}(a, b)$ indicating a normal distribution with mean a and variance b. The formulation is not more complicated for any alternative Gaussian noise model. Moreover, we have observed the condition state of N_s components. By using index $j = i - N_y$ ranging from 1 to N_s, component j is at location $\mathbf{x}_{N_y+j}$, where intensity is $f_{N_y+j} = f(\mathbf{x}_{N_y+j})$, the uncertain capacity of the component is $c_j \sim \mathcal{N}(\mu_{c,j}, \sigma_{c,j}^2)$, and its binary state S_j is -1 if the local intensity is above the capacity, and the component fails, and 1 if, otherwise, the component survives.

Also, we are interested in N_* locations, where there can be additional uninspected components: $\mathbf{x}_i^*$, $f_i^* = f(\mathbf{x}_i^*)$ and S_i^* indicate the location, demand and state of uninspected component i, while its uncertain capacity is $c_i^* \sim \mathcal{N}(\mu_{c,i}^*, \sigma_{c,i}^{*2})$. Each of the capacities of inspected and uninspected components is assumed to be independent of all other variables.

Given vectors of observations $\mathbf{y} = [y_1 \cdots y_{N_y}]^T$ and $\mathbf{s} = [S_1 \cdots S_{N_s}]^T$, prior functions μ_f and k_f, noise variance σ_ε^2, and capacity parameters $\{\mu_{c,1}, \ldots, \mu_{c,N_s}\}$, $\{\sigma_{c,1}, \ldots, \sigma_{c,N_s}\}$, $\{\mu_{c,1}^*, \ldots, \mu_{c,N_*}^*\}$, and $\{\sigma_{c,1}^*, \ldots, \sigma_{c,N_*}^*\}$, our goal is to approximate posterior distributions $p(\mathbf{f}_*|\mathbf{s}, \mathbf{y})$ and $p(\mathbf{s}_*|\mathbf{s}, \mathbf{y})$, of vectors listing unmeasured intensity $\mathbf{f}_* = [f_1^* \cdots f_{N_*}^*]^T$ and unobserved components' states $\mathbf{s}_* = [S_1^* \cdots S_{N_*}^*]^T$.

3. Inference Process

3.1 Sequential Bayesian Updating: Processing Observations of Intensity and of Components' Condition

For performing inference, we assume the probabilistic graphical model in Fig. 1: vectors $\mathbf{f}_y = [f_1 \cdots f_{N_y}]^T$ and $\mathbf{f}_s = [f_{1+N_y} \cdots f_{N_s+N_y}]^T$ list demands in locations where intensity and component state is observed, respectively.

At the core of the graph, the three vectors $\mathbf{f}_y$, $\mathbf{f}_s$ and $\mathbf{f}_*$ are statistically interdependent. We do not assume any specific causal relation among these vectors, and we assign directions to the arrows following the easiness of presentation. The graph factorizes the vectors' joint distribution as the product of $p(\mathbf{f}_y)$, $p(\mathbf{f}_s|\mathbf{f}_y)$ and $p(\mathbf{f}_*|\mathbf{f}_s, \mathbf{f}_y)$, the product of the latter two factors being

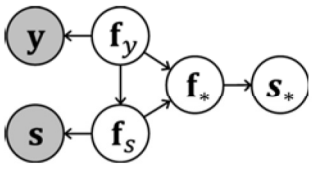

Fig. 1. Probabilistic Graphical Model for the Inference Process

$p(\mathbf{f}_*, \mathbf{f}_s|\mathbf{f}_y)$.

We begin by processing observations $\mathbf{y}$. First we update the distribution of $\mathbf{f}_y$ using Bayes' formula:

$$p(\mathbf{f}_y|\mathbf{y}) \propto p(\mathbf{f}_y)p(\mathbf{y}|\mathbf{f}_y) \tag{2}$$

then we predict vectors $\mathbf{f}_*$ and $\mathbf{f}_s$:

$$p(\mathbf{f}_*, \mathbf{f}_S|\mathbf{y}) = \int p(\mathbf{f}_y|\mathbf{y})p(\mathbf{f}_*, \mathbf{f}_s|\mathbf{f}_y)d\mathbf{f}_y \tag{3}$$

This formula is valid because, according to Fig. 1, $\{\mathbf{f}_*, \mathbf{f}_s\}$ is independent of $\mathbf{y}$, conditional on $\mathbf{f}_y$. We derive two additional distributions from this outcome: the posterior distribution of $\mathbf{f}_s$: $p(\mathbf{f}_s|\mathbf{y}) = \int p(\mathbf{f}_*, \mathbf{f}_s|\mathbf{y})d\mathbf{f}_*$ and the posterior conditional distribution of $\mathbf{f}_*$ given $\mathbf{f}_s$: $p(\mathbf{f}_*|\mathbf{f}_s, \mathbf{y}) = p(\mathbf{f}_*, \mathbf{f}_s|\mathbf{y})/p(\mathbf{f}_s|\mathbf{y})$. Then we process observations $\mathbf{s}$, again using Bayes' formula:

$$p(\mathbf{f}_s|\mathbf{y}, \mathbf{s}) \propto p(\mathbf{f}_s|\mathbf{y})P(\mathbf{s}|\mathbf{f}_s) \tag{4}$$

and we propagate the inference to $\mathbf{f}_*$:

$$p(\mathbf{f}_*|\mathbf{y}, \mathbf{s}) = \int p(\mathbf{f}_*|\mathbf{f}_s, \mathbf{y})p(\mathbf{f}_s|\mathbf{y}, \mathbf{s})d\mathbf{f}_s \tag{5}$$

Again, this is valid because $\mathbf{f}_*$ is independent of $\mathbf{s}$, conditional on $\mathbf{f}_s$. Finally, we get posterior distribution of components' state:

$$P(\mathbf{s}_*|\mathbf{s}, \mathbf{y}) = \int p(\mathbf{f}_*|\mathbf{y}, \mathbf{s})P(\mathbf{s}_*|\mathbf{f}_*)d\mathbf{f}_* \tag{6}$$

3.2 Processing Intensity Measures: Gaussian Processes for Regression

We can perform inference in Eqs. (2)-(3) exactly, adopting the scheme of Gaussian processes for regression. The prior distributions of $\mathbf{f}_y$, of augmented field vector $\mathbf{f}_+ = [\mathbf{f}_*^T \ \mathbf{f}_s^T]^T$ and of noise vector $\boldsymbol{\varepsilon} = [\varepsilon_1 \cdots \varepsilon_N]^T$ are:

$$\mathbf{f}_y \sim N(\boldsymbol{\mu}_{F_y}, \boldsymbol{\Sigma}_{F_y}); \ \mathbf{f}_+ \sim \mathcal{N}(\boldsymbol{\mu}_{F_+}, \boldsymbol{\Sigma}_{F_+}); \ \boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_e) \tag{7}$$

where mean vectors $\boldsymbol{\mu}_{F_y}$ and $\boldsymbol{\mu}_{F_+}$ are derived from the mean function μ_f, covariance matrixes $\boldsymbol{\Sigma}_{F_y}$ and $\boldsymbol{\Sigma}_{F_+}$ from covariance function k_f, while noise covariance matrix $\boldsymbol{\Sigma}_e$ is $\sigma_e^2\mathbf{I}$, following the assumption of white homoscedastic noise. Moreover, we let $\boldsymbol{\Sigma}_{F_+F_y}$ indicate the covariance matrix between $\mathbf{f}_y$ and $\mathbf{f}_+$, again derived from covariance function k_f.

We can process in closed form any observations $\mathbf{y}$ linearly related to $\mathbf{f}_y$ (Malings and Pozzi, 2016b). Here, we assume the simple model outlined in previous section: $\mathbf{y} = \mathbf{f}_y + \boldsymbol{\varepsilon}$. Hence, the marginal distribution of $\mathbf{y}$ is Gaussian, with mean vector $\boldsymbol{\mu}_Y = \boldsymbol{\mu}_{F_y}$ and covariance matrix $\boldsymbol{\Sigma}_Y = \boldsymbol{\Sigma}_{F_y} + \boldsymbol{\Sigma}_e$. The distribution of $\mathbf{f}_+$, posterior with respect to $\mathbf{y}$, is again normal, with mean vector and covariance matrix given by:

$$\boldsymbol{\mu}_{F_+|\mathbf{y}} = \boldsymbol{\mu}_{F_+} + \boldsymbol{\Sigma}_{F_+F_y}\boldsymbol{\Sigma}_Y^{-1}(\mathbf{y} - \boldsymbol{\mu}_Y); \ \boldsymbol{\Sigma}_{F_+|Y} = \boldsymbol{\Sigma}_{F_+} - \boldsymbol{\Sigma}_{F_+F_y}\boldsymbol{\Sigma}_Y^{-1}\boldsymbol{\Sigma}_{F_+F_y}^T \tag{8}$$

3.3 Processing Condition Observations: Gaussian Processes for Classification

Now that the measures in vector $\mathbf{y}$ have been processed, we can focus on inferring $\mathbf{f}_*$ from $\mathbf{s}$. By partitioning posterior mean vector $\boldsymbol{\mu}_{F_+|\mathbf{y}}$ and covariance matrix $\boldsymbol{\Sigma}_{F_+|Y}$ of $\mathbf{f}_+$, as computed in Eq. (8), we get corresponding quantities for $\mathbf{f}_*$ and $\mathbf{f}_s$. Let $\boldsymbol{\mu}_{F_*|\mathbf{y}}$ and $\boldsymbol{\mu}_{F_S|\mathbf{y}}$ indicate the posterior mean vectors for $\mathbf{f}_*$ and $\mathbf{f}_s$, $\boldsymbol{\Sigma}_{F_*|\mathbf{y}}$ and $\boldsymbol{\Sigma}_{F_S|\mathbf{y}}$ the corresponding covariance matrixes, and let $\boldsymbol{\Sigma}_{F_*F_S|\mathbf{y}}$ indicate the posterior covariance matrix of $\mathbf{f}_*$ and $\mathbf{f}_s$.

We perform the inference process in the domain of vector $\mathbf{f}_s$. When processing observations $\mathbf{s}$, the Gaussian distribution $p(\mathbf{f}_s|\mathbf{y})$, posterior respect to processing $\mathbf{y}$, plays the role of a prior distribution. The likelihood function of component j in $\mathbf{f}_s$ is:

$$P(S_j|f_{j'}) = \Phi(a_j) \tag{9}$$

where $j' = j + N_y$, $a_j = S_j\frac{\mu_{c,j} - f_{j'}}{\sigma_{c,j}}$ and Φ is the cumulative standard normal distribution. The joint log-likelihood, related to the observation of all conditions, is:

$$\log P(\mathbf{s}|\mathbf{f}_s) = \sum_{j=1}^{N_S} \log P(S_j|f_{j'}) \tag{10}$$

By applying Eq. (4) we get the updated distribution of $\mathbf{f}_s$. However, the likelihood in Eq. (9) is not conjugate to the Gaussian prior, and no closed form is available for the posterior distribution. We can adopt deterministic numerical methods for approximating the posterior distribution with a Gaussian distribution. The Laplace's method and the Expectation Propagation method are alternative approaches for identifying this distribution.

3.3.1 Laplace's Method

The Laplace's Method (LM) identifies the value $\hat{\mathbf{f}}_s$ that maximizes the product of prior distribution and likelihood function, as this product is proportional to the posterior distribution, according to Eq. (4). The logarithm Ψ of this product is:

$$\Psi(\mathbf{f}_s) = \log P(\mathbf{s}|\mathbf{f}_s) + \log p(\mathbf{f}_s|\mathbf{y}) \tag{11}$$

By substituting Eqs. (9)-(10) and the Gaussian prior in Eq. (11), we get the following form for the gradient of Ψ:

$$\nabla\Psi(\mathbf{f}_s) = -\boldsymbol{\Sigma}_{F_S|\mathbf{y}}^{-1}(\mathbf{f}_s - \boldsymbol{\mu}_{F_S|\mathbf{y}}) + \boldsymbol{\alpha} \tag{12}$$

where $\boldsymbol{\alpha} = \nabla \log P(\mathbf{s}|\mathbf{f}_s)$ is a vector of N_s components, and its component j is:

$$\alpha_j = \frac{\partial \log P(S_j|f_{j'})}{\partial f_{j'}} = -\frac{S_j}{\sigma_{c,j}}\frac{\varphi(a_j)}{\Phi(a_j)} \tag{13}$$

where φ is the standard normal distribution density. The Hessian matrix of Ψ is:

$$\nabla\nabla\Psi(\mathbf{f}_s) = -\boldsymbol{\Sigma}_{F_S|\mathbf{y}}^{-1} - \overline{\boldsymbol{\Sigma}}^{-1} \tag{14}$$

where $\overline{\boldsymbol{\Sigma}}^{-1} = -\nabla\nabla \log P(\mathbf{s}|\mathbf{f}_s)$ is a diagonal matrix, and the entry j on its diagonal is:

$$b_j = -\frac{\partial^2 \log P(S_j|f_{j'})}{(\partial f_{j'})^2} = \frac{1}{\sigma_{c,j}^2}\frac{\varphi(a_j)}{\Phi(a_j)}\left[a_j + \frac{\varphi(a_j)}{\Phi(a_j)}\right] \tag{15}$$

Function Ψ is concave, and the Newton-Raphson method (Theodoridis, 2015) can be applied to iteratively identify value $\hat{\mathbf{f}}_S$, relying on the gradient and the Hessian matrix defined in Eqs. (12), (14). The posterior probability is then approximated with a Gaussian distribution with mean $\boldsymbol{\mu}_{F_S|\mathbf{s},y}$ equal to $\hat{\mathbf{f}}_s$, and with covariance matrix $\boldsymbol{\Sigma}_{F_S|\mathbf{s},\mathbf{y}}$ related to the Hessian matrix computed at $\hat{\mathbf{f}}_s$:

$$p(\mathbf{f}_s|\mathbf{s},\mathbf{y}) \cong \mathcal{N}(\mathbf{f}_s, \hat{\mathbf{f}}_s, [-\nabla\nabla\Psi(\hat{\mathbf{f}}_s)]^{-1}) = \mathcal{N}(\mathbf{f}_s, \boldsymbol{\mu}_{F_S|s,y}, \boldsymbol{\Sigma}_{F_S|s,y}) \quad (16)$$

We propagate the inference to vector $\mathbf{f}_*$, as in Eq. (5), by following the same scheme of Eq.8. The resulting posterior distribution $p(\mathbf{f}_*|\mathbf{y},\mathbf{s})$ is Gaussian, with mean vector:

$$\boldsymbol{\mu}_{F_*|s,y} = \boldsymbol{\mu}_{F_*|y} + \boldsymbol{\Sigma}_{F_*F_S|\mathbf{y}} \boldsymbol{\Sigma}_{F_S|\mathbf{y}}^{-1} (\boldsymbol{\mu}_{F_S|s,y} - \boldsymbol{\mu}_{F_S|y}) \quad (17)$$

and covariance matrix:

$$\boldsymbol{\Sigma}_{F_*|\mathbf{s},\mathbf{y}} = \boldsymbol{\Sigma}_{F_*|\mathbf{y}} - \boldsymbol{\Sigma}_{F_*F_S|\mathbf{y}} [\boldsymbol{\Sigma}_{F_S|\mathbf{y}} + \overline{\boldsymbol{\Sigma}}]^{-1} \boldsymbol{\Sigma}_{F_*F_S|\mathbf{y}}^T \quad (18)$$

with $\overline{\boldsymbol{\Sigma}}$ computed at $\hat{\mathbf{f}}_S$. By following this scheme, we obtain an approximate Gaussian posterior distribution of the intensity for all locations of interest, which can be used for risk assessment.

3.3.2 Expectation Propagation

Expectation Propagation (EP) is an alternative approach for approximating the posterior probability with a Gaussian distribution. While LM is based on a local investigation of the posterior distribution around its maximum, EP aims at minimizing the KL divergence between the actual and the approximate distribution. The likelihood function related to the observed condition at location $\mathbf{x}_j$ is approximately proportional to a Gaussian density with mean $\tilde{\mu}_j$ and standard deviation $\tilde{\sigma}_j$:

$$P(S_j|f_j) \tilde{\propto} \mathcal{N}(f_j, \tilde{\mu}_j, \tilde{\sigma}_j^2) \quad (19)$$

as if intensity $\tilde{\mu}_j$ was measured there, with noise variance $\tilde{\sigma}_j^2$. Once appropriate values of mean and variance are assigned to each location where the component state is observed, the posterior distribution of the demand in those locations is jointly Gaussian, with covariance matrix given by:

$$\boldsymbol{\Sigma}_{F_S|\mathbf{s},\mathbf{y}} = [\boldsymbol{\Sigma}_{F_S|\mathbf{y}}^{-1} + \tilde{\boldsymbol{\Sigma}}^{-1}]^{-1} \quad (20)$$

where diagonal matrix $\tilde{\boldsymbol{\Sigma}}$ lists terms $\{\tilde{\sigma}_1^2, \tilde{\sigma}_2^2, \ldots, \tilde{\sigma}_{N_S}^2\}$ on its diagonal, and with mean vector:

$$\boldsymbol{\mu}_{F_S|\mathbf{s},\mathbf{y}} = \boldsymbol{\Sigma}_{F_S|\mathbf{s},\mathbf{y}} [\boldsymbol{\Sigma}_{F_S|\mathbf{y}}^{-1} \boldsymbol{\mu}_{F_S|\mathbf{y}} + \tilde{\boldsymbol{\Sigma}}^{-1} \tilde{\boldsymbol{\mu}}] \quad (21)$$

where vector $\tilde{\boldsymbol{\mu}}$ lists terms $\{\tilde{\mu}_1, \tilde{\mu}_2, \ldots, \tilde{\mu}_{N_S}\}$. To identify the appropriate values of pair $\{\tilde{\mu}_j, \tilde{\sigma}_j^2\}$, we adapt the iterative method presented by Rasmussen and Williams (2006), as reported in the Appendix. Then, we use Eq. (17)-(18), with $\tilde{\boldsymbol{\Sigma}} = \overline{\boldsymbol{\Sigma}}$.

3.3.3 Monte Carlo estimation

Monte Carlo (MC) methods provide an alternative numerical approach for approximate inference. Using uniform sampling, we generate n_r samples of the field values in all the N_s locations where condition observations are collected, from Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}_{F_S|\mathbf{y}}, \boldsymbol{\Sigma}_{F_S|\mathbf{y}})$, posterior with respect to the intensity observations $\mathbf{y}$. Then, we assign to every sample k, $\mathbf{f}_{s(k)}$, weight $w_{(k)} = P(\mathbf{s}|\mathbf{f}_{s(k)})$, computed according to Eq. (10), and we normalize the n_r weights so that their sum is one: $\overline{w}_{(k)} = w_{(k)} / \Sigma_{l=1}^{n_r} w_{(l)}$. Eventually, we represent the posterior distribution, following the approach of Eq. (5), as:

$$p(\mathbf{f}_*|\mathbf{y},\mathbf{s}) \cong \Sigma_{k=1}^{n_r} \overline{w}_{(k)} p(\mathbf{f}_*|\mathbf{f}_{s(k)},\mathbf{y}) \quad (22)$$

For every sample k, $p(\mathbf{f}_*|\mathbf{f}_{s(k)},\mathbf{y})$ is a Gaussian distribution, whose variance is the same for all samples, and whose mean is an affine function of $\mathbf{f}_{s(k)}$, following the approach of Eq.8. In facts, that distribution represents the posterior probability after observing exactly the intensities listed in vector $\mathbf{f}_{s(k)}$, in the N_s locations. While Eq. (22) represents the posterior distribution as a mixture of Gaussians, we can also easily get samples of $\mathbf{f}_*$ if needed, e.g. by generating samples in the joint domain of $\{\mathbf{f}_s, \mathbf{f}_*\}$.

To reduce the variance of the estimator, we can adopt the alternative MC approach of Importance Sampling (IS), selecting a proposal distribution h, on the domain of $\mathbf{f}_s$, that concentrates the samples where the posterior distribution is high. The weight associated to sample k is now:

$$w_{(k)} = \frac{h(\mathbf{f}_{s(k)}) p(\mathbf{s}|\mathbf{f}_{s(k)})}{p(\mathbf{f}_{s(k)}|\mathbf{y})} \quad (23)$$

After normalization of the weights, we can use again Eq. (22) for predicting the stressor intensity. In Eq. (23), the density at the denominator is Gaussian, and so is h if we select the outcome of LM or of EP as a proposal distribution. We conclude by noting that the estimate based on Monte Carlo, either with uniform or importance sampling, is consistent but not necessarily unbiased, because of the sample normalization related to Bayes' formula (MacKay, 2003).

3.3.4 Condition Assessment of Uninspected Components

Following LM or EP, the posterior distribution of $\mathbf{f}_*$ is approximated by a single Gaussian distribution. To implement Eq. (6) and, specifically, to obtain the marginal probability of failure of the uninspected components, we let $\mu_{\omega,i}^*$ and $\sigma_{\omega,i}^{*2}$ indicate the posterior mean and variance of the intensity at location $\mathbf{x}_i^*$. The corresponding posterior reliability index $\beta_{\omega,i}^*$ is (Der Kiureghian, 2005):

$$\beta_{\omega,i}^* = \frac{\mu_{c,i}^* - \mu_{\omega,i}^*}{\sqrt{\sigma_{c,i}^{*2} + \sigma_{\omega,i}^{*2}}} \quad (24)$$

and, consistently with Eq.9, the posterior probability of failure is $P_{F,\omega,i}^* = \Phi(-\beta_{\omega,i}^*)$.

When using MC methods, the posterior distribution is represented by a mixture of Gaussians, and the corresponding distribution of the states of uninspected components is:

$$P(\mathbf{s}_*|\mathbf{s},\mathbf{y}) \cong \Sigma_{k=1}^{n_r} \overline{w}_{(k)} \int p(\mathbf{f}_*|\mathbf{f}_{s(k)},\mathbf{y}) P(\mathbf{s}_*|\mathbf{f}_*) d\mathbf{f}_* \quad (25)$$

As distribution $p(\mathbf{f}_*|\mathbf{f}_{s(k)},\mathbf{y})$ is Gaussian, we obtain, for each sample, a corresponding failure probability for the component in location $\mathbf{x}_i^*$, following the approach of Eq. (24) to solve the integral in Eq. (25). Then, the posterior probability of failure is estimated using the weighted average of Eq. (25).

3.3.5 Extension to Systems with Multi-state Components

While our framework refers to binary components, we briefly outline how to extend it for dealing with multi-state components. The first issue is related to the processing of observations related to intermediate damage states. By observing that a component is

in an intermediate state, we conclude that the local field cannot be too high or too low as, in those cases, the component would be respectively undamaged or failed. A path to extend LM to multi-state classification is presented by Williams and Baber (1998) and reviewed by Rasmussen and Williams (2006) as "multi-class Laplace approximation". For EP, the likelihood functions related to these intermediate states may be closer to a Gaussian density, and so the approximation of Eq.19 may be even more acceptable. A recent work in this direction is that of Villacampa-Calvo and Hernández-Lobato (2017).

A second issue is the prediction of the condition of uninspected components among more than two possible states, once the inference of the latent Gaussian process has been performed. In the application of the next Section, we will illustrate how this can be done in a given setting.

4. Applications to Spatially Distributed Systems

4.1 Illustrative Application to a 1-D Domain

To illustrate the inference process, we define the following problem. In a one-dimensional domain, where the spatial coordinate x ranges from 0 to 10 Km, spatial stressor f is defined by a stationary Gaussian process, with mean $\mu_f(x) = 1.4$ and standard deviation $\sigma_f(x) = 0.5$ for every location x. The corresponding 95% credible region for the stressor, $CR_{95\%\pi}(f)$, is reported in each graph in the middle row of Fig. 2(a). The covariance function follows the squared exponential model:

$$k_f(x,x') = \sigma_f(x)\sigma_f(x')\exp\left[-\frac{(x-x')^2}{2\lambda^2}\right] \tag{26}$$

with correlation length $\lambda = 2$ Km. By using this model, we assume a smooth field, with correlation equal to 97%, 88%, 60% and 13% for a pair of points at distance 0.5 km, 1 km, 2 km and 4 km respectively. We observe the state of 30 components distributed along the domain: 28 are undamaged, while 2 failed at x equal to 6.12 km and 6.74 km respectively. The upper graphs in Fig. 2(a) show the observations in the domain: letters U and F indicate an Undamaged or a Failed condition (corresponding to $s = 1$ or $s = -1$) respectively. All components have independent capacities defined by standard deviation $\sigma_c = 0.3$, and mean μ_c. We consider different values of μ_c, to show how the results are affected by this parameter. In the left column in Fig. 2(a), $\mu_c = 2.146$, so that the prior probability of failure for each component is $P_{F,\pi} = 10\%$. The 95% credible region of the capacities is reported in the middle graph, as $CR_{95\%}(c)$. To cover the domain, we define a regular grid of $N_* = 1,000$ points, and the covariance related to points in the grid, or between them and observed locations, follows Eq. (26). That graph also reports the posterior credible region for the stressor in the grid, $CR_{95\%\omega}(f)$, as a function of the spatial coordinate, for different inference approaches: LM stands for Laplace's Method, EP for Expectation Propagation, MC for uniform Monte Carlo, IS for Importance Sampling Monte Carlo, where the proposal distribution h is the Gaussian approximation of the posterior distribution identified by LM. Both MC and IS estimates are based on 10^5 samples.

By observing the middle graph in the first column of Fig. 2(a), we note that the updated marginal stressor is not much affected by the observations of the undamaged components. For example, all 15 components for $x < 6$ km are undamaged, but the updated intensity is not significantly lower than the prior one: this happens because, as the prior probability of components being undamaged is high (it is $1 - P_{F,\pi} = 90\%$), the observation that many components survived is not much informative. On the contrary, the failure of two components indicates that the stressor could be high in a region, and the posterior field is significantly updated there: as the (prior) lower bound of $CR_{95\%}(c)$ is 1.55, it is unlikely that a stressor below that value could have destroyed any component, and so the lower bound of $CR_{95\%\omega}(f)$ is now above that level near the location of failures. All methods agree, but LM assigns a higher upper bound to $CR_{95\%\omega}(f)$ for low values of x. The graph below reports the corresponding prediction of failure or survival of a test component, depending on its location, in terms of posterior reliability index. Here we assume that $\mu_c^* = \mu_c$ and $\sigma_c^* = \sigma_c$, so that the prior probability of failure of this test component is also $P_{F,\pi}^* = P_{F,\pi} = 10\%$, corresponding to $\beta_\pi^* = 1.28$. This index grows in the regions where undamaged components are observed, and decreases where failures are observed. We note that, even if the impact of the updating to the intensity field looks minor for $x < 4$ km, the consequences in terms of reliability index are significant, as $CR_{95\%}(c)$ overlaps much less with $CR_{95\%\omega}(f)$ than with $CR_{95\%\pi}(f)$ in that region. Finally, we note that the EP inference is in good agreement with that performed by MC and IS, and so we conclude that EP provides accurate results.

We note that no method performs any explicit inference about the capacity. The capacity of the observed component is "marginalized out" in the definition of the likelihood function in Eq. (9). If needed, we can obtain the posterior distribution of the capacity of observed components, and we would tend to infer, for example, that the posterior expected capacity of the failed components is lower than the prior expected one. However, there is no need to update the distribution of the capacity of the uninspected components, as it has been assumed to be marginally independent of those of the observed ones. Hence, when we predict the fate of the uninspected components, the inference of the stressor intensity is the only relevant one.

In the second column, all assumptions are the same as before (including the set of binary data collected in-field), except for the higher components' capacity, that is now assumed to be $\mu_c = 2.756$, so that $P_{F,\pi}^* = 1\%$, corresponding to $\beta_\pi^* = 2.33$. The results are similar as before, except for the following features. As the components are now stronger, the observation of some failures makes the updated stressor model go higher. Also, the impact of the observations of undamaged components is less significant, both in terms of intensity and reliability index, as we do expect (with 99% probability) that components do not fail. The third column refers to an even higher capacity: $\mu_c = 3.202$, so that $P_{F,\pi}^* = 0.1\%$, and graphs in Fig. 2(b) to $\mu_c = 3.568$, 3.887,

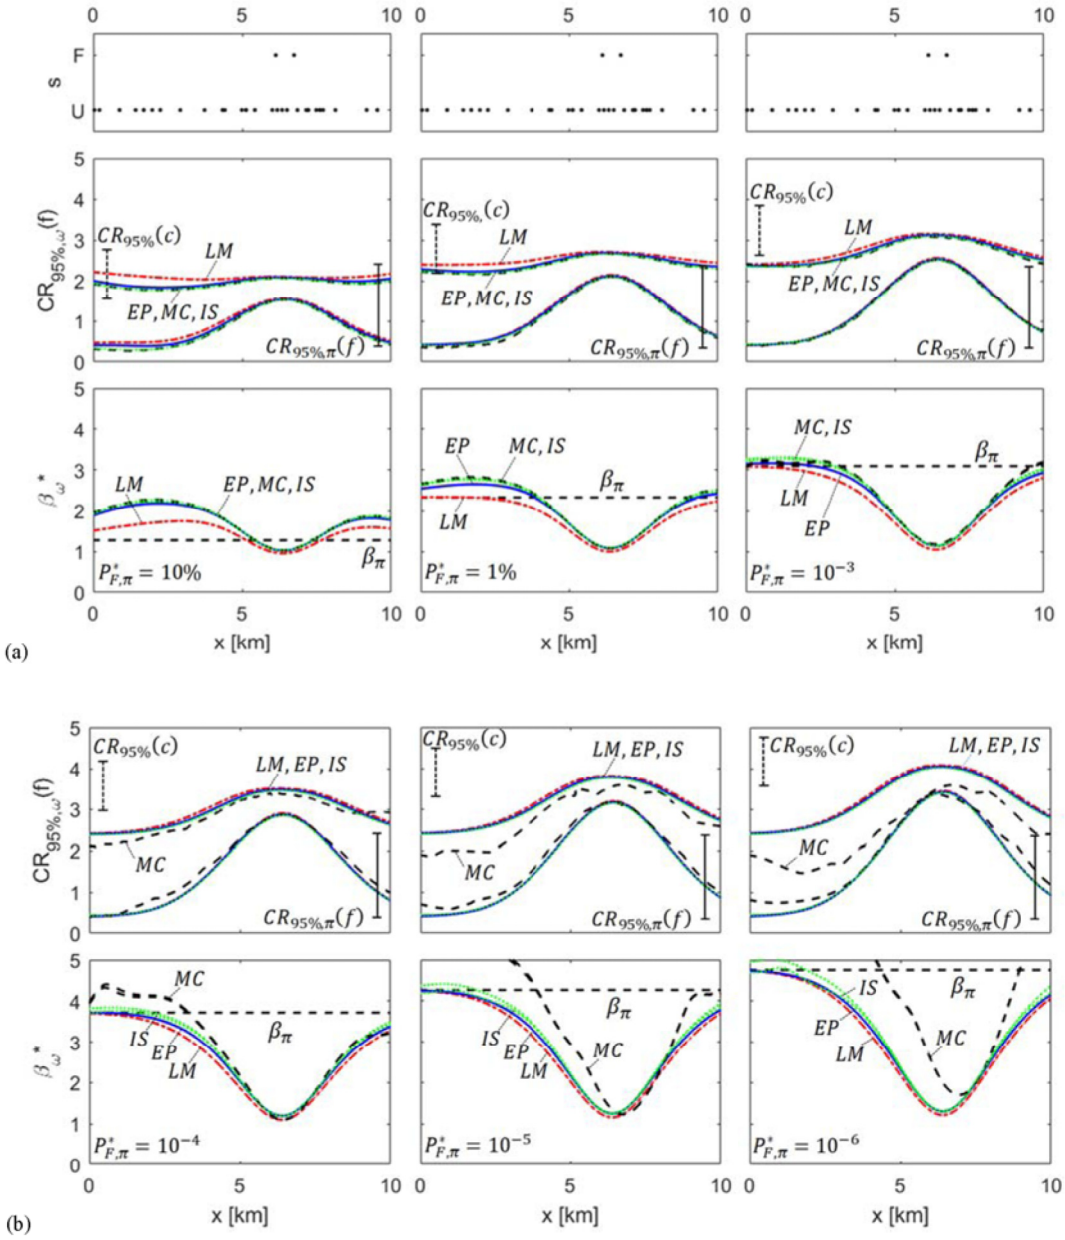

Fig. 2. (a) Inference Process in the 1-D Domain. Binary Observations (upper graphs), Posterior Credible Region of the Stressor (middle graphs) and Posterior Reliability Index (lower graphs), for Components' Prior Failure Probability from 10% to 0.1%, (b) Corresponding Graphs for Components' Prior Failure Probability from 10^{-4} to 10^{-6}

4.171 respectively, so that $P^*_{F,\pi} = 10^{-4}$, 10^{-5}, 10^{-6}. We observe that the same trend persists: observations of failure of stronger components make the updated field go higher, while observations of stronger components being undamaged are less significant. Moreover, we note that the inferences of IS, EP and LM are in good agreement (specifically, the agreement of LM with IS and EP is even better than that for weaker components).

However, as the considered components are stronger and less prone to failure, the MC estimate becomes significantly biased. The effective number of samples can be defined as:

$$n_{eff} = 1/\sum_{k=1}^{n_r} \overline{w}_{(k)}^2 \tag{27}$$

and the ratio of effective samples as $r_{eff} = n_{eff}/n_r$. Table 1 reports n_{eff} and r_{eff}, for IS and MC, depending on the components' capacity. For high capacity, $P_{F,\pi}$ is low, and so is the probability of observing two failures. Hence, most simulated fields are in poor agreement with the observations. When $P_{F,\pi} = 10^{-6}$, n_{eff} goes down to 2.3 samples for MC, so that the posterior belief is represented just by a pair of samples, and this explains the poor performance shown in Fig. 2. n_{eff} is always high for IS, with the corresponding r_{eff} even increasing from 58% to 92%: this may

Table 1. Number and Ratio of Effective Samples in the MC and IS Simulations

$P_{F,\pi}$	$n_{eff,IS}$	$n_{eff,MC}$	$r_{eff,IS}$	$r_{eff,MC}$
10%	58,274	9,253	58%	9.3%
1%	76,748	3,078	77%	3.1%
10^{-3}	84,428	464	84%	4.6×10^{-3}
10^{-4}	88,347	41	88%	4.1×10^{-4}
10^{-5}	90,769	6.2	91%	6.2×10^{-5}
10^{-6}	92,455	2.3	92%	2.3×10^{-5}

indicate that the Gaussian distribution identified by LM for approximating the posterior is more accurate when $P_{F,\pi}$ is low, as also indicated, in Fig. 2, by the better agreement with EP.

In the previous analysis we used $\mu_c = \mu_c^*$, but this is not at all a requirement of any method outlined above. Generally, the moments of each capacity (for observed or unobserved components) can be different. For example, we may observe the condition of strong components and predict the behavior of weak ones.

In Fig. 3, we expand previous example considering two levels of components' damage: low (L) and high (H), defined by a probability of occurrence $P_{F,\pi}$ equal to 1% and to 10^{-4} respectively. In Fig. 3(a), we assume that the observed failures refer to any damage level (L or H), while the lower and upper sets of lines refer to the prediction of any damage and to the prediction of high damage respectively. Hence, the lower set is identical to that in the second column of Fig. 2(a). We note that the observations that no damage (not even low damage) occurred in the region where $x < 6$ km significantly increases the reliability index related to high damage in that region. Also, both deterministic methods perform poorly in that region. In Fig. 3(b), we assume that the observed failures refer to high damage, while the "undamaged" components may actually be at no damage or low damage. The inference process is as before, and the upper set of lines are identical to the first column of Fig. 2(a). The lower set of lines refers to the prediction of any damage level: we conclude that it is highly likely that components around location 6.5 km are damaged, when two highly damaged components have been observed in the surrounding region.

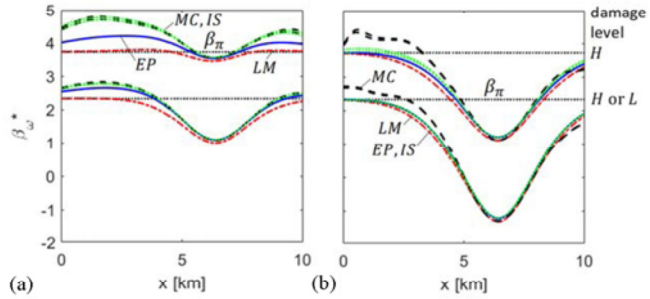

Fig. 3. Prediction for Two Damage Levels, in the 1-D Domain Application: Observations of Damage Refer to: (a) Any Damage or to (b) High Damage Only

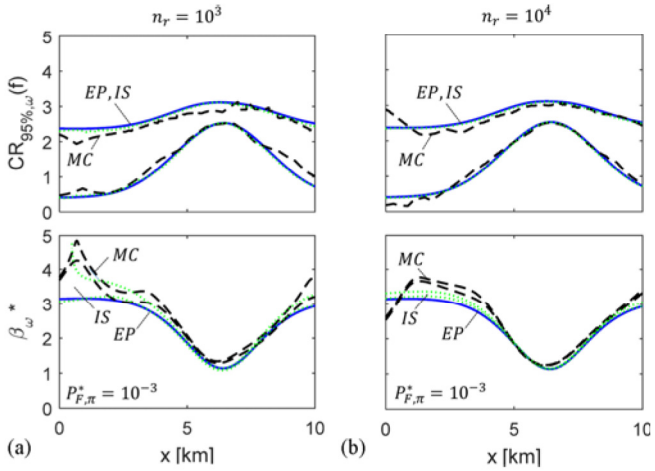

Fig. 4. Monte Carlo Estimates for Different Numbers of Samples: for: (a) 1,000 Samples and for (b) 10,000 Samples

Figure 4 shows the impact of the number of samples, n_r, in the reliability assessment based on MC and IS. We refer to the case when $\mu_c = 3.202$, so that $P_{F,\pi}^* = 0.1\%$, as reported in the third column of Fig.2(a). While, as reported above, that graph made use of 10^5 samples, Fig. 4(a) and (b) uses 10^3 and 10^4 samples, respectively. As expected, with a lower number of samples the accuracy degrades, and the estimates can be significantly biased. The number of samples needed for a good estimate depends on a combination of factors: how rare the evidence collected by the observations is, and how low the probability of the event we want to estimate is.

4.2 Application to Seismic Risk Assessment in a 2-D Domain

We now apply the approach to the risk assessment of a set of assets exposed to a seismic event. Fig. 4 shows the region where the assets are located, as described by two spatial coordinates (x_1 and x_2). We assume that an event of magnitude 7.5 occurs at location {40 km, 5 km} and that the Peak Spectral Acceleration (PSA) for a natural period of 0.55s is the appropriate intensity measure. We use the Akkar and Bommer (2010) attenuation law for modeling the seismic intensity, assuming a normal fault type and rock soil class. As the conditional distribution of PSA is lognormal, we define the Gaussian field f as the natural logarithm of the PSA in ms^{-2}. From the attenuation law, we derive the intra-event and the inter-event terms: the first term defines a common uncertainty affecting the entire field, while we adopt a squared exponential correlation function (as in the Eq. (26)) for the second term, with correlation length $\lambda = 7$ km. So, by combining these two terms, we get the covariance function k_f.

We consider two types of components. Their capacities are defined by independent log-normal random variables, with standard deviation of the log-capacity $\sigma_c = 0.3$ (approximately corresponding to a 30% coefficient of variation in the actual capacity) and expected capacity depending on the components' type: the expected log-capacity for strong components is

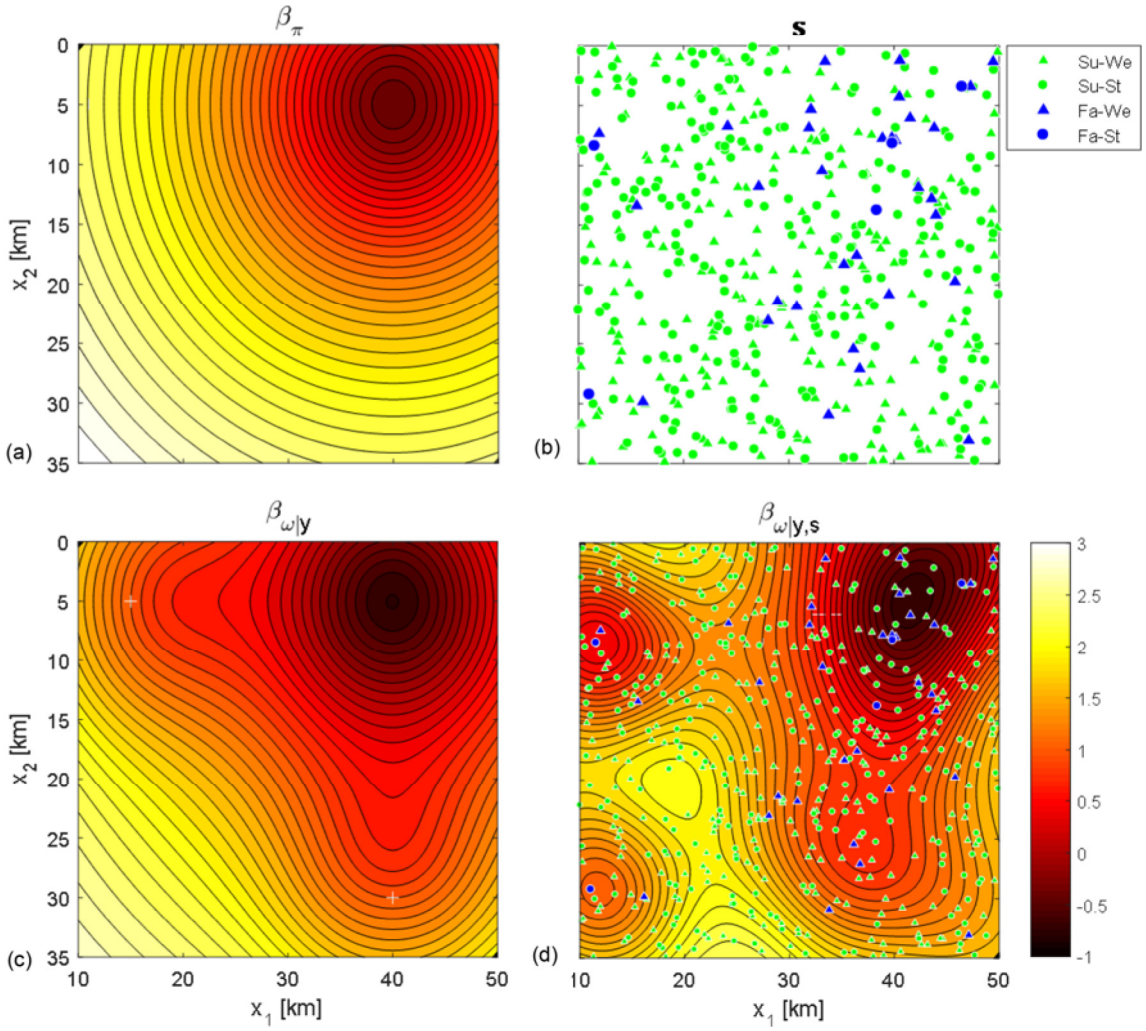

Fig. 5. (a) Prior Map of Reliability Index, (b) Locations and Conditions of the Observed Components, (c) Posterior Map of Reliability Index, after Processing PSA Measures and (d) after Also Processing Observations of Components' Condition

$\mu_c = 2.3$, and for weak components is $\mu_c = 1.6$, corresponding to a median actual capacity of about 10 ms^{-2} and about 5 ms^{-2}, respectively. Fig. 5(a) shows the map of the prior reliability index for a weak component, depending on its location: it ranges from -0.37, above the epicenter, where the median PSA is 5.86 ms^{-2}, up to 3.1. The map is plotted on an 81 × 71 grid, so that $N_* = 5,751$. We observe the condition of 500 components in the region: among the 250 strong ones, 5 fail, and among the 250 weak ones, 32 fail. Their position and observed condition are reported in Fig. 5(b), where "Su-We" indicates a survived weak component, "Su-St" a survived strong one, "Fa-We" a failed weak one and "Fa-St" a failed strong one. The log-PSA has been measured in two locations, with precision defined by $\sigma_\varepsilon = 10^{-4}$, at coordinates {15 km, 5 km} and {40 km, 30 km}. Both measures (y_1 and y_2) are equal to 1.3, corresponding to measuring a PSA equal to 3.67 ms^{-2}, significantly higher than the median prior value in both locations, of 1.92 ms^{-2}. The locations of these measures are marked by white crosses in Fig. 5(c), and the map also shows the updated reliability index for weak components.

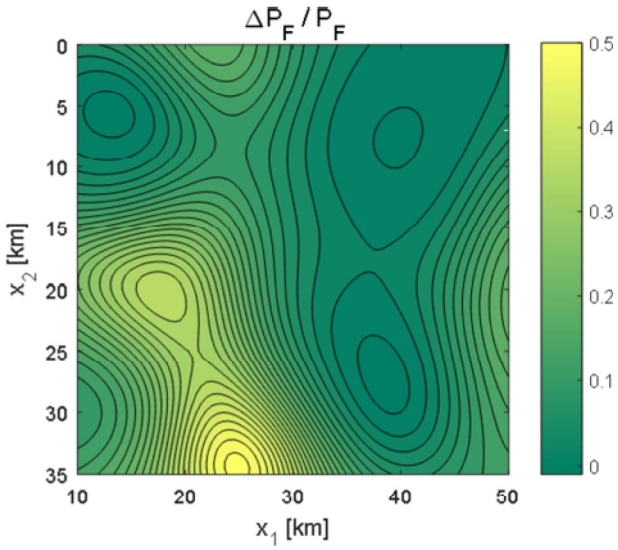

Fig. 6. Comparison between Reliability Assessment by LM and EP Methods

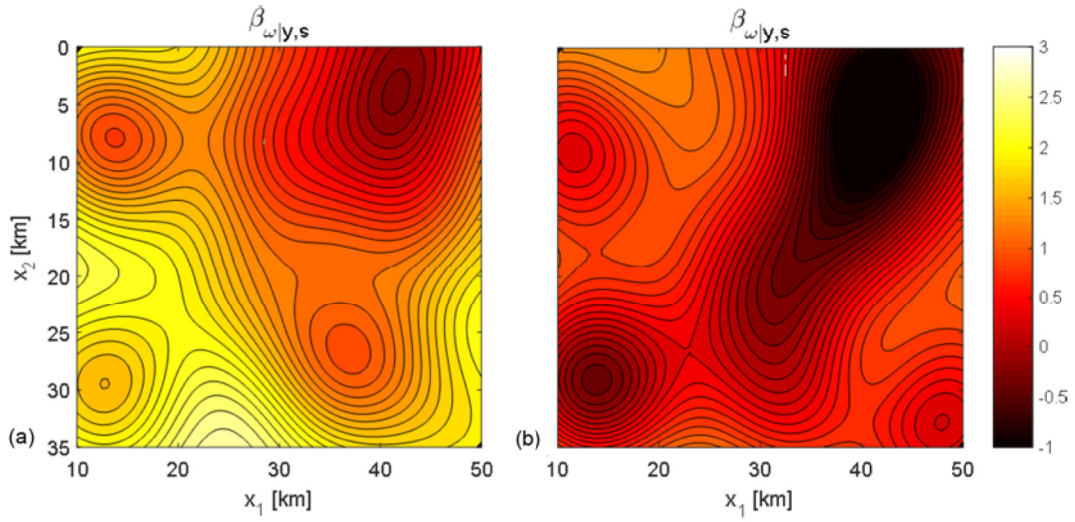

Fig. 7. Posterior Map of Reliability Index Assuming that All Observed Components were: (a) Weak and that All were, (b) Strong

Fig. 5(d) shows the posterior reliability index, after processing the observations on components' condition: this map fuses the information contained in (b) and (c). As expected, reliability grows near survived components, and decreases near failed components. The outcomes refer to the EP algorithm.

To compare the assessment based on the EP and LM methods, we define ΔP_F as the difference between the probability of failure obtained by LM and that obtained by EP, and we plot with difference, normalized by the estimate obtained by EM, in Fig. 6. The most significant difference is at location {25 km, 35 km}, where LM assigns a 1.28% probability of failure, while EP assigns 0.85%, so that the normalized difference is 0.51.

Figure 7 illustrates the role played in the inference process by the assumptions on the components' capacity. Keeping the components' location as in Fig. 5(b), Fig. 7(a) reports the posterior map of reliability index (as above, for a weak component) assuming all observed components were weak, and Fig. 7(b) the corresponding map (again, for a weak component) assuming all observed components were strong. We note how observing the failure of strong components has a higher impact in decreasing the assessed reliability. In the area above the epicenter, the minimum reliability index in the latter analysis is -1.4, corresponding to a probability of failure of about 92% (but we show all values below -1 at the value of -1, for preserving the same color scale of other maps).

We close this session with a remark on computational complexity. For this application, both methods use about 5-6 iterations (for the Newton-Raphson optimization in LM and for identifying the approximate likelihood functions in EM, respectively). On a laptop, with 16Gb of RAM and Intel® Core™ i7-7Y75 CPU @1.30GHz processor, the processing of components' states takes 0.2s for LM and 3s for EP, using MathWorks Matlab. The prediction for all points in the grid takes about 0.6s, for each of the two methods.

5. Conclusions

The approach outlined above allows for inferring, probabilistically, the condition of uninspected components based on the detected states of others. The setting closely resembles that of the classification problem in supervised learning where, after having observed a set of labeled points, we train a classifier to predict the class of unobserved locations in the features' space. Traditional methods for classification capture the intuition that components may tend to fail in certain sub-regions (e.g. using parametric functions, as with logistic regression) or that components may tend to fail in cluster (e.g. using non-parametric functions, as with k-nearest neighbor classification). However, much relevant information cannot be included in that traditional analysis. On the contrary, the framework outlined in this paper is based on a latent field that models an engineering quantity of interest, i.e. the stressor intensity; this quantity is probabilistically related to the components' condition, but it can also be directly measured. The prior distribution of this field comes from engineering models, and the framework allows for consistently processing measures of the stressor and of the components' state. When we adopt the Gaussian model, we can easily handle hundreds of measures, for projecting the results to thousands of locations, spending just few seconds (or even less than one second, using LM) on a common laptop.

A few points deserve further investigation. One issue is related to the applicability of the framework to specific engineering problems. In the case of seismic risk assessment, demands as PSA can be modeled, after the logarithm transformation, as a Gaussian field. Similarly, the capacity of some components can be modeled with Gaussian variables, after the same transformation, as in the illustrative example. However, if the significant natural period is not uniform among components, more complicate correlation models, as that proposed by Loth and Baker (2013), should be used. It is still an open question how a full recording of

ground motion in one location should globally affect the inference about the PSA in other locations, for all natural periods. Another remark is related to the accuracy of the estimate based on deterministic approximate methods, as EP and LM. As the posterior distribution is not Gaussian, we cannot expect the reliability assessment to be exact. Neither of these methods is specifically focused on the accuracy of predicting the failure of any component: LM is based on a local approximation at the mode of the posterior, while EP on the minimization of the KL divergence. In the numerical examples we have investigated, we have not noted significant errors, except for the case where the reliability index is above 4. We have compared those methods with MC and IS estimates. Uniform MC is ineffective when the collected measures are highly informative, either due to the high number of observations or to the observation of some rare events, because, in these cases, most generated samples are in poor agreement with the evidence and, consequently, they are insignificant. The deterministic methods can provide a suitable proposal distribution, for reducing the variance in the MC estimate, using IS. Further variance reduction can be achieved by moving the proposal distribution towards higher values of the stressor, to better cover the "design point" for unobserved components, as attempted by Pozzi & Der Kiureghian (2013). However, it is generally hard to define a proposal distribution able to cover all these design points in a satisfactory way.

Acknowledgements

We acknowledge the support of NSF project CMMI #1653716, titled "CAREER: Infrastructure Management under Model Uncertainty: Adaptive Sequential Learning and Decision Making".

References

Akkar, S. and Bommer, J. J. (2010). "Empirical equations for the prediction of pga, pgv, and spectral accelerations in europe, the mediterranean region, and the Middle East." *Seismological Research Letters,* Vol. 81, No. 2, pp. 195-206, DOI: 10.1785/gssrl.81.2.195.

Bensi, M., Der Kiureghian, A., and Straub, D. (2014). "Framework for post-earthquake risk assessment and decision making for infrastructure systems." *ASCE-ASME J. Risk Uncertainty Eng. Syst., Part A: Civ. Eng.,* 04014003-1:17, DOI: 10.1061/AJRUA6.0000810.

Cavalieri, F., Franchin, P., Gehl, P., and D'Ayala, D. (2017). "Bayesian networks and infrastructure systems: Computational and methodological challenges." *In Risk and Reliability Analysis: Theory and Applications* (pp. 385-415). Springer International Publishing, DOI: 10.1007/978-3-319-52425-2_17.

Der Kiureghian, A. (2005). "First-and second-order reliability methods." *In Singhal S, Ghiocel DM, Nikolaidis E (eds) Engineering design reliability handbook.* Ch.14-1, DOI: 10.1201/9780203483930.ch14

Krige, D. G. (1951). "A statistical approach to some basic mine valuation problems on the Witwatersrand." *J. of the Chem., Metal. and Mining Soc. of South Africa*, Vol. 52, No. 6, pp. 119-139, DOI: 10520/AJA0038223X_4792.

Loth, C. and Baker, J. W. (2013). "A spatial cross-correlation model of spectral accelerations at multiple periods." *Earthquake Engng. Struct. Dyn.*, Vol. 42, No. 3, pp. 397-417, DOI: 10.1002/eqe.2212.

MacKay, D. J. C. (2003). *Information Theory, Inference, and Learning Algorithms*, Cambridge University Press, Cambridge, United Kingdom.

Malings, C. and Pozzi, M. (2016a). "Conditional entropy and value of information metrics for optimal sensing in infrastructure systems." *Structural Safety*, Vol. 60, pp. 77-90, DOI: 10.1016/j.strusafe.2015.10.003.

Malings, C. and Pozzi, M. (2016b). "Value of Information for Spatially Distributed Systems: application to sensor placement." *Reliability Engineering & System Safety,* Vol. 154, pp. 219-233, DOI: 10.1016/j.ress.2016.05.010.

Pozzi, M. and Der Kiureghian, A. (2013). "Gaussian Bayesian network for system reliability of bridges." in G. Deodatis; B.R. Ellingwood and D.M. Frangopol (Eds.) Safety, Reliability, Risk and Life-Cycle Performance of Structures and Infrastructures, *Proceedings of the 11th International Conference on Structural Safety And Reliability (ICOSSAR2013)*, New York, NY. 16-20 June 2013.

Rasmussen, C. E. and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning*, MIT Press, Masachusetts, United States.

Theodoridis, S. (2015). *Machine learning a bayesian and optimization perspective,* Academic Press, Masachusetts, United States.

Villacampa-Calvo, C. and Hernández-Lobato, D. (2017). *Scalable multi-class Gaussian process classification using expectation propagation,* arXiv preprint arXiv:1706.07258.

Williams, C. K. I. and Barber, D. (1998) "Bayesian classification with Gaussian processes." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 20, No. 12, pp. 1342-1351, DOI: 10.1109/34.735807.

Yue, Y., Pozzi, M., Zonta, D., and Zandonini, R. (2012). "Bayesian networks for post-earthquake assessment of bridges." in "Bridge Maintenance, Safety, Management, Resilience and Sustainability." Leiden (NL): CRC Press/Balkema, 2012. *Proceedings of 6th International Conference on Bridge Maintenance, Safety and Management (IABMAS 2012), Stresa, 8-12 Jul, 2012.*

Appendix: Details on EP

In this section we outline how to identify parameters $\{\tilde{\mu}_1, \tilde{\mu}_2, \ldots, \tilde{\mu}_{N_S}\}$ and $\{\tilde{\sigma}_1^2, \tilde{\sigma}_2^2, \ldots, \tilde{\sigma}_{N_S}^2\}$ to fill vector $\tilde{\boldsymbol{\mu}}$ and matrix $\tilde{\boldsymbol{\Sigma}}$ in Eqs. (20)-(21). We do this iteratively, by initializing all $\tilde{\mu}_j$ and $\tilde{\sigma}_j^{-2}$ to zero, and identifying each pair $\{\tilde{\mu}_j, \tilde{\sigma}_j^2\}$ assuming the other values as known. Following closely Rasmussen and Williams (2006), the "cavity" distribution q_{-j}, marginal posterior for $f_{j'}$ including likelihood functions for all components except component j itself, is $q_{-j}(f_{j'}) = \mathcal{N}(f_{j'}, \mu_{-j}, \sigma_{-j}^2)$. Its parameters are:

$$\sigma_{-j}^2 = (\sigma_j^{-2} - \tilde{\sigma}_j^{-2})^{-1}\ ;\ \mu_{-j} = \sigma_{-j}^2(\sigma_j^{-2}\mu_j - \tilde{\sigma}_j^{-2}\tilde{\mu}_j) \tag{28}$$

where μ_j and σ_j^2 are entry j on vector $\boldsymbol{\mu}_{F_S|y}$ and on the diagonal of matrix $\boldsymbol{\Sigma}_{F_S|y}$ respectively.

The product of the cavity distribution and the likelihood function can be approximated by a Gaussian distribution by moment matching, to minimize the KL divergence.

$$q_{-j}(f_{j'})P(S_j|f_{j'}) \propto \mathcal{N}(f_{j'}, \hat{\mu}_i, \hat{\sigma}_i^2) \propto \mathcal{N}(f_{j'}, \mu_{-j}, \sigma_{-j}^2)\mathcal{N}(f_{j'}, \tilde{\mu}_j, \tilde{\sigma}_j^2) \tag{29}$$

To identify moments $\hat{\mu}_i$ and $\hat{\sigma}_i^2$, we have to combine the parameters of components' capacity to that cavity distribution.

Generalizing the formulas in Rasmussen & Williams (2006), we start normalizing the parameters of the cavity distribution, obtaining:

$$\overline{\sigma}_{-j}^2 = \sigma_{-j}^2 \sigma_{c,j}^{-2};\ \overline{\mu}_{-j} = (\mu_{c,j} - \mu_{-j})/\sigma_{c,j} \tag{30}$$

Then, again from Rasmussen & Williams (2006), we compute the moments $\{\breve{\mu}_j,\ \breve{\sigma}_j^2\}$ in the normalized space of f_j, as

$$\breve{\sigma}_j^2 = \overline{\sigma}_{-j}^2 - \overline{\sigma}_{-j}^4 r_j \chi_j^2 (z_j + r_j)\ ;\ \breve{\mu}_j = \overline{\mu}_{-j} + s_j \overline{\sigma}_{-j}^2 r_j \chi_j \tag{31}$$

where $r_j = \varphi(z_j)/\Phi(z_j)$, $z_j = s_j \mu_{-j} \chi_j$ and $\chi_j = (1 + \overline{\sigma}_{-j}^2)^{-1/2}$. Then we get the posterior moments for the actual stressor by applying the inverse transformation respect to Eq. (30), obtaining:

$$\hat{\sigma}_i^2 = \breve{\sigma}_i^2 \sigma_{c,j}^2 \quad \hat{\mu}_i = \sigma_{c,j} \breve{\mu}_i + \mu_{c,j} \tag{32}$$

Finally, we identify the parameters of the approximate likelihood function by applying the inverse transformation respect to Eq. (28), obtaining:

$$\tilde{\sigma}_j^2 = (\hat{\sigma}_j^{-2} - \sigma_{-j}^{-2})^{-1} \quad \tilde{\mu}_j = \tilde{\sigma}_j^2 (\hat{\sigma}_j^{-2} \hat{\mu}_j - \sigma_{-j}^{-2} \mu_{-j}) \tag{33}$$

We iterate the updating of $\{\tilde{\mu}_j, \tilde{\sigma}_j^2\}$ for all observed components, and then again and again until we get convergence.