Health monitoring of civil structures: A MCMC approach based on a multi-fidelity deep neural network surrogate[†]

Matteo Torzoni [1,2,*] [ID], **Andrea Manzoni** [2] [ID] and **Stefano Mariani** [1] [ID]

1   Dipartimento di Ingegneria Civile ed Ambientale, Politecnico di Milano, Piazza L. da Vinci 32, 20133 Milano, Italy
2   MOX, Dipartimento di Matematica, Politecnico di Milano, Piazza L. da Vinci 32, 20133 Milano, Italy
*   Correspondence: matteo.torzoni@polimi.it
†   Presented at the 1st Online Conference on Algorithms, 27 September–10 October 2021.

Abstract: To meet the need for reliable real-time monitoring of civil structures, safety control and optimization of maintenance operations, this paper presents a computational method for the stochastic estimation of the degradation of the load bearing structural properties. Exploiting a Bayesian framework, the procedure sequentially updates the posterior probability of the damage parameters used to describe the aforementioned degradation, conditioned on noisy sensors observations, by means of Markov chain Monte Carlo (MCMC) sampling algorithms. To enable the analysis to run in real-time or close to, the numerical model of the structure is replaced with a data-driven surrogate used to evaluate the conditional likelihood. The proposed surrogate model relies on a multi-fidelity (MF) deep neural network (DNN), mapping the damage and operational parameters onto approximated sensor recordings. The MF-DNN is shown to effectively leverage information between multiple datasets, by learning the correlations across models with different fidelities without any prior assumption, ultimately alleviating the computational burden of the supervised training stage. The low fidelity (LF) responses are approximated by relying on proper orthogonal decomposition for the sake of dimensionality reduction, and a fully connected DNN. The high fidelity signals, that feed the MCMC within the outer-loop optimization, are instead generated by enriching the LF approximations through a deep long short-term memory network. Results relevant to a specific case study demonstrate the capability of the proposed procedure to estimate the distribution of damage parameters, and prove the effectiveness of the MF scheme in outperforming a single-fidelity based method.

Keywords: structural health monitoring; Markov chain Monte Carlo; deep learning; multi-fidelity; reduced order modeling; damage identification.

1. Introduction

Civil structures and infrastructures are critical for the life of the world population and play a strategic role for the global economy [1]. Aging and ever-increasing extreme loading conditions threaten existing and new structural systems, stressing the need of real-time structural health monitoring (SHM) procedures to detect and identify any deviation from the damage-free baseline [2].

Vibration-based SHM techniques investigate the structural health by recording and analyzing the vibration response (e.g., acceleration or displacement multivariate time series) of the monitored structure. Two competitive SHM approaches can be formally distinguished [3]: the model-based one, e.g., [4,5], and the data-based one, e.g., [6,7]. The former is usually implemented through an updating strategy of a physics-based model on the basis of measured experimental data, which attempts to estimate the location and the extent of the occurred structural changes. On the other hand, data-driven methods are based on a Machine Learning (ML) paradigm that, once trained, can be used as a black-box tool. ML systems automatically learn how the features, originated from the

recorded data, are statistically correlated with the sought damage patterns [8]. After the advent of Deep Learning (DL) [9], that can incorporate the selection and extraction of optimized features into the end-to-end learning processes, the feature engineering stage has been progressively automatized.

This work proposes a non-parametric output-only approach to the damage localization problem (see for instance [10,11]), leveraging a synergic combination of Multi-Fidelity (MF) data-driven meta-modeling and Bayesian parameter identification. The probability distribution of the unknown damage parameters is approximated through a Markov chain Monte Carlo (MCMC) sampling algorithm.

MCMC has been applied in Bayesian model updating and model class selection in structural mechanics (see, e.g., [12]) as well as in SHM (see, e.g., [13]). In this work, MCMC is used to construct a Markov chain of the sought damage parameters, whose limit distribution is the target probability distribution. The probability distribution is sequentially updated by exploring the support of the damage parameters with a density of steps proportional to the unknown posterior distribution. The sampling acceptance is governed by the evidence of the current parameters to represent sparse dynamic response measurements, as provided by a sensors network, by means of a data-driven surrogate model.

Because handling FE simulations within a MCMC analysis is computationally impractical, a FE model capable of simulating the effect of damage on the structural response is adopted only to build labelled datasets of vibration recordings for known damage positions (see for instance [14]). A data-driven surrogate model is instead adopted to map operational and damage parameters to the associated vibration signals in place of the FE model. Such surrogate is based on a multi-fidelity deep neural network (MF-DNN) trained on synthetic data of multiple fidelities, a ML paradigm adopted and extended for instance in [15,16]. Specifically, a limited amount of high fidelity (HF) data and a lot of cheaper low fidelity (LF) data are considered. This type of meta-modeling is useful to alleviate the high demand during training of, potentially expensive to collect, HF data. Indeed, the LF data supply useful information on the trends of HF data, allowing the MF-DNN to enhance the prediction accuracy only leveraging few HF data in comparison to the single-fidelity method [17].

2. SHM Methodology

The proposed methodology is detailed as follows. The composition of the datasets used to train the surrogate model is specified in Sec. 2.1, the numerical models behind these datasets are discussed in Sec. 2.2, the MF-DNN surrogate model is described in Sec. 2.3, and the setup of the MCMC analysis used for the sake of damage localization is explained in Sec. 2.4.

2.1. Datasets definition

The LF and HF datasets, respectively $\mathbf{D}_{\text{LF}}$ and $\mathbf{D}_{\text{HF}}$, are built from the assembly of I_{LF} and I_{HF} instances, as follows

$$\mathbf{D}_{\text{LF}} = \{(\mathbf{x}_i^{\text{LF}}, \mathbf{U}_i^{\text{LF}})\}_{i=1}^{I_{\text{LF}}}, \quad \mathbf{D}_{\text{HF}} = \{(\mathbf{x}_j^{\text{HF}}, \mathbf{U}_j^{\text{HF}})\}_{j=1}^{I_{\text{HF}}}. \tag{1}$$

Each LF instance is provided by a LF model of the structure to be monitored in undamaged conditions, and consists of the input parameters $\mathbf{x}_i^{\text{LF}} \in \mathbb{R}^{\text{N}_{\text{par}}^{\text{LF}}}$ defining the operational conditions, i.e. the loadings acting on the structure during the i-th instance, and the relative LF vibration recordings $\mathbf{U}_i^{\text{LF}}(\mathbf{x}_i^{\text{LF}}) = [\mathbf{u}_1^{\text{LF}}, \ldots, \mathbf{u}_{N_u}^{\text{LF}}]_i \in \mathbb{R}^{N_u \times L}$ shaped as N_u arrays of length L. The HF counterpart is provided by a HF model of the same structure, which also accounts for the presence of structural damage and internal damping. Each HF instance consists of the input parameters $\mathbf{x}_j^{\text{HF}} \in \mathbb{R}^{\text{N}_{\text{par}}^{\text{HF}}}$, defining the operational and damage conditions, with $\text{N}_{\text{par}}^{\text{HF}} > \text{N}_{\text{par}}^{\text{LF}}$, and the associated HF vibration recordings $\mathbf{U}_j^{\text{HF}}(\mathbf{x}_j^{\text{HF}}) \in \mathbb{R}^{N_u \times L}$. The structural damage is modeled as a selective reduction of the

84 material stiffness, applied to a subdomain identified by the spatial coordinates of its
85 center $\boldsymbol{\theta}_j \subset \mathbf{x}_j^{\text{HF}}$. For simplicity, the same sampling frequency and monitored degrees of
86 freedom (dofs) are considered for the two fidelities, but there are no restrictions on this
87 respect. Each instance refers to a time window $(0, T)$, short enough to assume frozen
88 operational, environmental, and damage conditions. In the reminder of the paper the
89 indexes i, j will be dropped.

2.2. Datasets population

The monitored structure is modeled as an elastic continuum discretized in space by
means of a FE triangulation. The HF numerical model results from the semi-discretized
form of the elasto-dynamic problem defined over the FE mesh. On the other hand, in
order to ease the construction of a large LF dataset, a projection-based model order
reduction strategy for parametrized systems is adopted to build the LF model, see
e.g. [10]. To this aim, the reduced basis method [18] relying on the Proper Orthogonal
Decomposition (POD)-Galerkin approach is considered. Hence, the LF approximation
is obtained as a linear combination of POD-basis functions, yet not accounting for the
presence of damage and structural damping. The LF and HF numerical models read
respectively as

$$
\begin{cases}
\mathbf{M}^R \ddot{\mathbf{d}}^R(t) + \mathbf{K}^R \mathbf{d}^R(t) = \mathbf{f}^R(\mathbf{x}^{\text{LF}}), & t \in (0, T) \\
\mathbf{d}^R(0) = \mathbf{W}^\top \mathbf{d}_0 \\
\dot{\mathbf{d}}^R(0) = \mathbf{W}^\top \dot{\mathbf{d}}_0,
\end{cases}
\tag{2}
$$

$$
\begin{cases}
\mathbf{M}\ddot{\mathbf{d}}(t) + \mathbf{C}(\mathbf{x}^{\text{HF}}(\boldsymbol{\theta}))\dot{\mathbf{d}}(t) + \mathbf{K}(\mathbf{x}^{\text{HF}}(\boldsymbol{\theta}))\mathbf{d}(t) = \mathbf{f}(\mathbf{x}^{\text{HF}}), & t \in (0, T) \\
\mathbf{d}(0) = \mathbf{d}_0 \\
\dot{\mathbf{d}}(0) = \dot{\mathbf{d}}_0,
\end{cases}
\tag{3}
$$

91 where the superscripts L and H are omitted from all the arrays for simplicity, while
92 the superscript R stands for *reduced*. Having denoted by: $t \in (0, T)$ the time co-
93 ordinate; $\mathbf{d}(t) \in \mathbb{R}^M$, $\dot{\mathbf{d}}(t) \in \mathbb{R}^M$ and $\ddot{\mathbf{d}}(t) \in \mathbb{R}^M$ the vectors of nodal displace-
94 ments, velocities and accelerations, respectively, whereas M is the number of dofs;
95 $\mathbf{M} \in \mathbb{R}^{M \times M}$ the mass matrix; $\mathbf{C}(\mathbf{x}^{\text{HF}}(\boldsymbol{\theta})) \in \mathbb{R}^{M \times M}$ the damping matrix, modeled as
96 Rayleigh damping for mathematical convenience; $\mathbf{K}(\mathbf{x}^{\text{HF}}(\boldsymbol{\theta})) \in \mathbb{R}^{M \times M}$ the stiffness
97 matrix; $\mathbf{f}(\mathbf{x}^{\text{LF}}), \mathbf{f}(\mathbf{x}^{\text{HF}}) \in \mathbb{R}^M$ the vectors of nodal forces; $\mathbf{d}_0$ and $\dot{\mathbf{d}}_0$ the initial conditions
98 at $t = 0$; $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_{M_R}] \in \mathbb{R}^{M \times M_R}$ the matrix gathering the $M_R \ll M$ retained
99 POD-basis functions; $\mathbf{M}^R, \mathbf{K}^R, \mathbf{f}^R(\mathbf{x}^{\text{LF}}), \mathbf{d}^R(t)$ the reduced arrays, playing the same role
100 of the FE matrices but with dimension ruled by M_R instead of M. It has to be noted that,
101 even if in this case the two fidelities differ through the presence of structural damage and
102 viscous damping in the HF model, the proposed computational framework is general
103 and can be arbitrarily adapted to different modeling choices.

104 The datasets $\mathbf{D}_{\text{LF}}$ and $\mathbf{D}_{\text{HF}}$ are populated accordingly to Eq. (1) by sampling the
105 parametric input spaces, respectively defined by a uniform probability distribution over
106 $\mathbf{x}^{\text{LF}}$ and $\mathbf{x}^{\text{HF}}$, via latin hypercube sampling. The relevant vibration recordings $\mathbf{U}^{\text{LF}}$ and
107 $\mathbf{U}^{\text{HF}}$ are extracted from $\mathbf{d}^{\text{LF}}$ and $\mathbf{d}^{\text{HF}}$, respectively, through a Boolean operation.

2.3. MF-DNN surrogate model

The MF-DNN $\mathcal{NN}_{\text{MF}}$ is composed of a LF neural network $\mathcal{NN}_{\text{LF}}$, trained on low-
cost data, which is used as baseline model, and a HF neural network $\mathcal{NN}_{\text{HF}}$, trained on
few HF data, which is used to adaptively learn the correlation between LF and HF data.
The overall evaluation of $\mathcal{NN}_{\text{MF}}$ reads as

$$
\hat{\mathbf{U}}^{\text{HF}} = \mathcal{NN}_{\text{MF}}(\mathbf{x}^{\text{HF}}, \mathbf{x}^{\text{LF}}) = \mathcal{NN}_{\text{HF}}(\mathbf{x}^{\text{HF}}, \hat{\mathbf{U}}^{\text{LF}}), \quad \hat{\mathbf{U}}^{\text{LF}} = \texttt{reshape}[\mathbf{Y}(\frac{1}{\boldsymbol{\omega}} \odot \mathcal{NN}_{\text{LF}}(\mathbf{x}^{\text{LF}}))].
\tag{4}
$$

109 Here: $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_{M_{\text{LF}}}] \in \mathbb{R}^{L_{\text{concat}} \times M_{\text{LF}}}$, with $L_{\text{concat}} = L \times N_u$, is a matrix gathering
110 M_{LF} POD-basis functions built upon $\mathbf{D}_L$ and used to compress the LF data in order

111 to ease the complexity of $\mathcal{NN}_{\text{LF}}$; $\mathcal{NN}_{\text{LF}}$ is a fully connected DNN, mapping the LF
112 input parameters onto the POD-basis coefficients; $\omega \in \mathbb{R}^{M_{\text{LF}}}$ is a vector of numbers
113 linearly decreasing from 1 to 0.2, used to weight the regression over the POD-basis
114 coefficients by their relative importance; $\odot$ denotes the Hadamard product; the reshape
115 operation is used to recast the reconstructed LF signals from a single vector of size
116 L_{concat} into N_u arrays of length L; $\mathcal{NN}_{\text{HF}}$ is a long short-term memory (LSTM) NN that,
117 as more appropriate to solve time-dependent problems, is adopted to map the HF input
118 parameters and the approximated LF signals onto the HF signals.

119 *2.4. Damage localization via MCMC*

Accordingly to the Bayes' rule, the posterior probability density function (pdf) of
the damage parameters θ, viewed as random variables, conditioned on the observed
signals $\mathbf{U}^{\text{EXP}}_{1,\dots,N_{\text{obs}}}$ is

$$p(\theta|\mathbf{U}^{\text{EXP}}_{1,\dots,N_{\text{obs}}}, \mathcal{NN}_{\text{MF}}) = \frac{p(\mathbf{U}^{\text{EXP}}_{1,\dots,N_{\text{obs}}}|\theta, \mathcal{NN}_{\text{MF}})p(\theta, \mathcal{NN}_{\text{MF}})}{\int p(\mathbf{U}^{\text{EXP}}_{1,\dots,N_{\text{obs}}}|\theta, \mathcal{NN}_{\text{MF}})p(\theta, \mathcal{NN}_{\text{MF}})\,d\theta}, \tag{5}$$

where: $p(\theta, \mathcal{NN}_{\text{MF}})$ is the prior of θ; $p(\mathbf{U}^{\text{EXP}}|\theta, \mathcal{NN}_{\text{MF}})$ is the likelihood of the evidence,
which measures the goodness of fit of $\mathcal{NN}_{\text{MF}}$ to $\mathbf{U}^{\text{EXP}}$ given the parameters θ. By
assuming that the uncertainties follow a Gaussian distribution, the likelihood function
can be assumed Gaussian too thanks to the central limit theorem:

$$p(\mathbf{U}^{\text{EXP}}_{1,\dots,N_{\text{obs}}}|\theta, \mathcal{NN}_{\text{MF}}) = \prod_{k}^{N_{\text{obs}}} \frac{1}{(\sqrt{2\pi})^L \sqrt{|\Sigma_c|}} \exp\left(-\frac{\frac{1}{L}\sum_{\tau=1}^{L}[(\mathbf{e}_\tau^\top \Delta_k)^\top \Sigma_c^{-1}(\mathbf{e}_\tau^\top \Delta_k)]}{2}\right). \tag{6}$$

120 Here: N_{obs} is the batch size of the processed observations; $\Delta_k = \mathbf{U}^{\text{EXP}}_k - \hat{\mathbf{U}}^{\text{HF}}(\mathbf{x}^{\text{HF}}(\theta), \mathbf{x}^{\text{LF}})$
121 is the prediction error relevant to the k-th observation, assumed independent between
122 different time instants and modeled as a Gaussian random vector with zero mean and
123 covariance matrix $\Sigma_c \in \mathbb{R}^{N_u \times N_u}$, which describes the spatial correlation of prediction
124 errors due to modeling errors and measurement noise (for further details see e.g. [19]);
125 $\mathbf{e}_\tau$ is a Boolean vector with a single non-zero entry in τ-th position, used as row extractor
126 for the relevant time step.

127 To avoid the expensive computation of the integral at the denominator of Eq. (5), a
128 MCMC sampling algorithm is adopted to approximate the posterior pdf
129 $p(\theta|\mathbf{U}^{\text{EXP}}_{1,\dots,N_{\text{obs}}}, \mathcal{NN}_{\text{MF}})$. Specifically, the posterior pdf is sequentially updated accord-
130 ingly to the Metropolis-Hastings (MH) algorithm [20]. The MH algorithm simulates
131 a chain of θ samples distributed according to the posterior, with each sample only de-
132 pending on the previous one. This generate a random walk in the space of θ, where
133 each point is hit with a frequency proportional to its probability. Hence, the stationary
134 distribution of the Markov chain, under the assumption of ergodicity, asymptotically
135 approaches the target pdf.

136 Let $q(\xi|\theta)$ be the considered *proposal* probability density function and
137 $\delta(\theta) = p(\mathbf{U}^{\text{EXP}}_{1,\dots,N_{\text{obs}}}|\theta, \mathcal{NN}_{\text{MF}})p(\theta, \mathcal{NN}_{\text{MF}})$ for the sake of simplicity. The MH algorithm
138 recursively simulate the next Markov chain sample θ_{k+1} from the current sample θ_k,
139 with $k = 1, \dots, L_{\text{chain}}$, as follows [21]: sample a candidate ξ from $q(\xi|\theta_k)$; compute the
140 ratio $\alpha = \frac{\delta(\xi)q(\theta_k|\xi)}{\delta(\theta_k)q(\xi|\theta_k)}$; accept the candidate ξ with probability $\min\{1, \alpha\}$ and store it as
141 next state of the chain, i.e. $\theta_{k+1} = \xi$, otherwise reject it and keep the current state of the
142 chain, i.e. $\theta_{k+1} = \theta_k$.

After L_{chain} states are evaluated, the burn-in period of the chain (i.e. the initial
transitory phase) is removed to eliminate the initialization effect. The resulting chain is
ultimately thinned up to $\tilde{L}_{\text{chain}} = \frac{L_{\text{chain}}}{k_T}$, with k_T a small fixed integer, in order to remove
dependencies among consecutive samples. In that way a sufficiently long stable Markov
chain is obtained. The target distribution can be approximated via histograms and the

posterior expected values and covariance can be approximated with the empirical mean and covariance of the $\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_{\tilde{L}_{\text{chain}}}$ samples:

$$\boldsymbol{\mu_\theta} = \mathbb{E}(\boldsymbol{\theta}|\mathbf{U}^{\text{EXP}}_{1,\ldots,\text{N}_{\text{obs}}}, \mathcal{NN}_{\text{MF}}) \approx \frac{1}{\tilde{L}_{\text{chain}}} \sum_{l=1}^{\tilde{L}_{\text{chain}}} \boldsymbol{\theta}_l \,, \tag{7}$$

$$\text{cov}(\boldsymbol{\theta}|\mathbf{U}^{\text{EXP}}_{1,\ldots,\text{N}_{\text{obs}}}, \mathcal{NN}_{\text{MF}}) \approx \frac{1}{\tilde{L}_{\text{chain}} - 1} \sum_{l=1}^{\tilde{L}_{\text{chain}}} [\boldsymbol{\theta}_l - \boldsymbol{\mu_\theta}][\boldsymbol{\theta}_l - \boldsymbol{\mu_\theta}]^\top \,. \tag{8}$$

3. Virtual experiment

The proposed method is validated on the digital twin shown in Fig. 1. The HF model in Eq. (3) is obtained from a FE discretization resulting in $M = 4659$ dofs and integrated in time using the Newmark method. The structure is made of concrete, whose mechanical properties are: Young's modulus $E = 30$ GPa; Poisson's ratio $\nu = 0.2$; density $\rho = 2500$ kg/m^3. The structure is excited at the tip by a distributed load $q(t)$, acting on an area of (0.3×0.3) m^2, as depicted in Fig. 1. The load $q(t)$ varies in time according to $q(t) = Q\sin(2\pi ft)$, where $Q \in [1,5]$ kPa and $f \in [10,60]$ Hz respectively denote the load amplitude and frequency, collected as $\mathbf{x}^{\text{LF}} = (Q, f)^\top$. Damage is introduced by reducing the material stiffness by 25% within the subdomain Ω, which is a box $(0.3 \times 0.3 \times 0.4)$ m^3 as depicted in Fig. 1. The target position of this reduction is given by the coordinates of its center and can be identified with a single abscissa $\theta_\Omega \in [0.15, 7.55]$ m running along the axis of the structure. Hence, the input parameters of the HF part are collected as $\mathbf{x}^{\text{HF}} = (Q, f, \theta_\Omega)^\top$. Also the Rayleigh damping matrix, which account for a 5% damping ratio on the first 4 structural modes, is affected by the damage through the stiffness matrix. Synthetic displacement recordings $\mathbf{u}_n(t)$, with $n = 1, \ldots, N_u$, are collected from $N_u = 8$ dofs at the bottom surface, mimicking a monitoring system arranged as depicted in Fig. 1. Each recording lasts for a time interval $(0, T = 1$ s$)$, providing $L = 200$ data points.
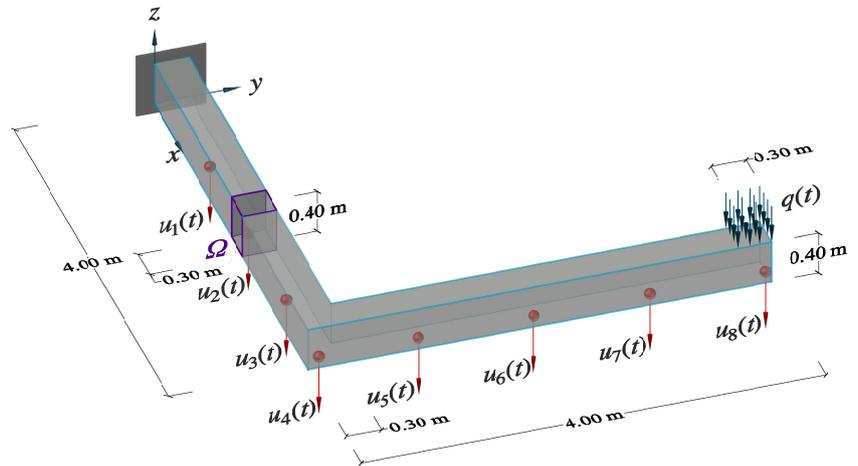
Figure 1. Physics-based digital twin of the monitored structure.

The reduced-order model in Eq. (2), i.e. the LF model used to construct $\mathbf{D}_{\text{LF}}$, has been built performing a POD upon 40.000 snapshots in time, collected while exploring the parametric input space $\mathbf{x}^{\text{LF}}$. 14 POD-bases are selected and stored in matrix $\mathbf{W}$, in place of the original 4659 dofs, after having fixed a suitable tolerance on the energy norm of the reconstruction error ($\text{tol}_{\text{POD}} = 10^{-3}$), for further details see, e.g., [10,14].

For the training of the surrogate model in Eq. (4), $I_{\text{LF}} = 10.000$ and $I_{\text{HF}} = 1.000$ instances have been collected from the LF and HF model, respectively. Concerning the compression of the LF data for the sake of prior dimensionality reduction, 104 POD-bases have been selected ($\text{tol}_{\text{POD}} = 10^{-3}$) and stored in matrix $\mathbf{Y}$, in place of the 1600 data points featured by each instance.

172   The mean squared error and the mean absolute error have been used as loss
173   functions for the training of $\mathcal{N}\mathcal{N}_{\text{LF}}$ and $\mathcal{N}\mathcal{N}_{\text{HF}}$, respectively, together with the Adam
174   optimization algorithm [22]. The implementation has been carried out through the
175   `Tensorflow`-based `Keras` API [23], running on an `Nvidia GeForce RTX 3080 GPU` card.
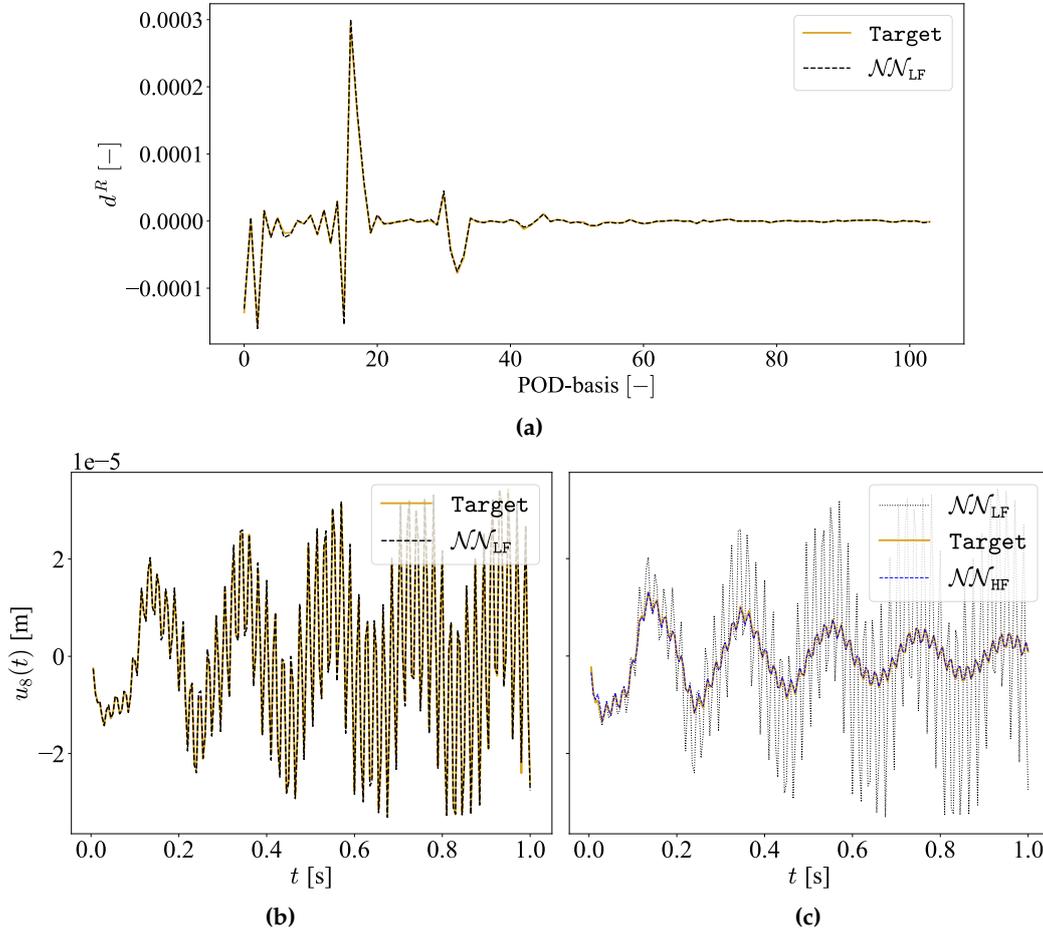
Figure 2. Reconstruction capacity of $\mathcal{N}\mathcal{N}_{\text{MF}}$: (a) regression over the POD-basis coefficients relative to a compressed LF signal; (b) decompressed LF signal; (c) regression over the HF signal.

176   An example of the reconstruction capabilities achieved by the surrogate model
177   is shown in Fig. 2 for the monitored gdl $u_8(t)$, where the outcome of the regression
178   over the POD-basis coefficients, ruled by the $\mathcal{N}\mathcal{N}_{\text{LF}}$, and the corresponding expanded
179   LF signal are reported together with the signal enrichment, provided by the $\mathcal{N}\mathcal{N}_{\text{HF}}$.
180   To quantify the accuracy of the predicted signals, the Pearson correlation coefficients
181   (PCC) between predicted and ground truth HF signals are adopted as a measure
182   of fitness. The PCC coefficients are evaluated with respect to 40 testing instances
183   generated with the HF model while exploring the parametric input space $\mathbf{x}^{\text{HF}}$. The
184   minimum PCC value over the 40 testing instances for each monitored channel is re-
185   spectively $\{0.983; 0.988; 0.994; 0.995; 0.998; 0.998; 0.998; 0.998\}$, which largely validate
186   the performance of the surrogate model. The other way around, if the $\mathcal{N}\mathcal{N}_{\text{HF}}$ is em-
187   ployed without being coupled with the $\mathcal{N}\mathcal{N}_{\text{LF}}$, the maximum PCC value drops to
188   $\{0.605; 0.603; 0.601; 0.601; 0.791; 0.735; 0.709; 0.696\}$, showing the utility of the MF setting
189   that outperforms the single-fidelity based method.
190   In the absence of experimental data, the Bayesian estimation of the damage param-
191   eter θ_Ω is simulated by considering pseudo-experimental instances, generated with the
192   HF model, that have been corrupted by adding independent, identically distributed
193   Gaussian noise, featuring a signal-to-noise ratio equal to 80 to each vibration recording.

194 Batches of $\mathtt{N_{obs}} = 3$ observations relative to the same damage condition but different
195 operational conditions are processed during the evaluation of the likelihood in Eq. (6).
196 The prior pdf $p(\theta_\Omega, \mathcal{NN}_{\mathtt{MF}})$ is taken as uniform, while, to account for the bounded
197 domain in which θ_Ω can fall, a truncated Gaussian centered on the last accepted state is
198 considered for the proposal $q(\boldsymbol{\xi}|\theta_\Omega)$. The adaptive Metropolis [24] algorithm is adopted
199 in order to ease the calibration of the proposal distribution, enabling its covariance to be
200 tuned on the basis of past samples as the sampling evolves. The MCMC algorithm is
201 run for 5000 samples, the first 500 of which are removed to get rid of the burn-in period.
202 The obtained chain is ultimately thinned by discarding 3 samples over 4 to remove
203 dependencies among consecutive samples.

204 Two examples of MCMC analyses are reported in Fig. 3, showing the generated
205 Markov chains alongside the estimated posterior mean and credibility intervals. In both
206 cases, the damage parameter θ_Ω, here normalized between 0 and 1, is properly identified.
207 It has to be noted that the larger uncertainty in the second case is somehow expected;
208 indeed, given the structural layout and the placing of the sensors, the sensitivity of
209 measures to damage positions far apart from the clamped side is smaller.
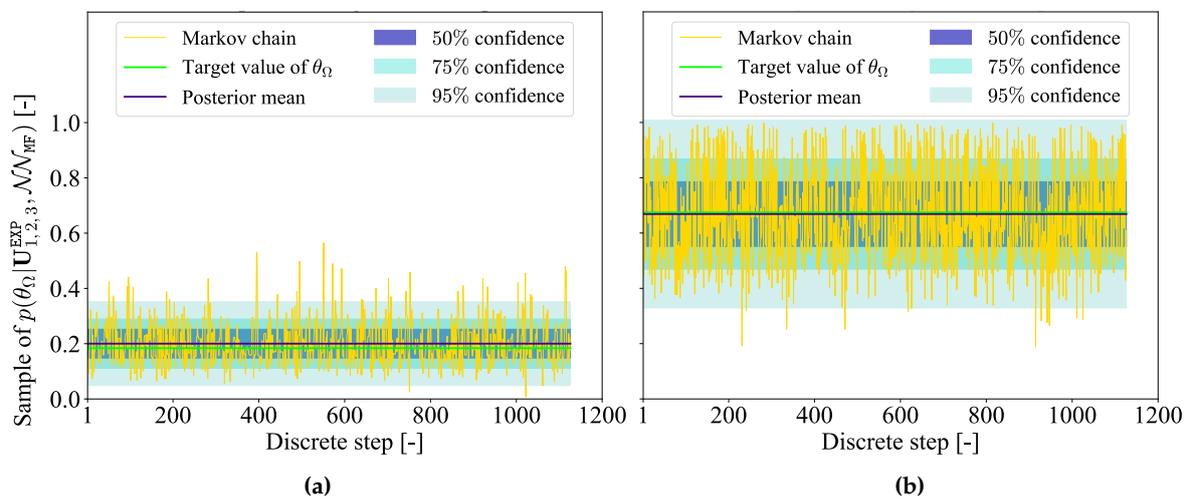
Figure 3. Examples of MCMC analyses.

4. Conclusions

211 This paper has presented a stochastic approach for SHM, here applied to the prob-
212 lem of damage localization in case of slow damage progression. The presence of damage
213 has been postulated as already detected, and only the localization task has been ana-
214 lyzed. The Bayesian identification of damage parameters is achieved through a MCMC
215 sampling algorithm, adopted to approximate their posterior distribution conditioned on
216 a set of measurements. Few investigations are present in literature involving the use of
217 MCMC for the health monitoring of civil structures, and this is the first one considering
218 a MF-DNN surrogate model to accelerate the computation of the conditional likelihood.
219 The surrogate model learns from simulated data of multiple fidelities, i.e. few HF data
220 and several inexpensive LF data, such to alleviate the computational burden of the
221 supervised training stage. The method has been assessed on a numerical case study,
222 showing remarkable accuracy and proving to be insensitive to the effect of measurement
223 noise and varying operational conditions.

224 Besides the need of validating the proposed methodology within a suitable exper-
225 imental setting, the next studies will extended the Bayesian identification also to the
226 parameters controlling the operational conditions. Moreover, a usage monitoring tool
227 powered by a suitable data-driven paradigm will be considered to provide useful prior
228 knowledge as opposite to an informative flat prior.

References

1. Flah, M.; Nunez, I.; Ben Chaabene, W.; Nehdi, M.L. Machine Learning Algorithms in Civil Structural Health Monitoring: A Systematic Review. *Archives of Computational Methods in Engineering* **2021**, *28*, 2621–2643. doi:10.1007/s11831-020-09471-9.

2. Mariani, S.; Azam, S.E. Health Monitoring of Flexible Structures Via Surface-mounted Microsensors: Network Optimization and Damage Detection. Proceedings of the 5th International Conference on Robotics and Automation Engineering, 2020, pp. 81–86. doi:10.1109/ICRAE50850.2020.9310827.

3. Worden, K. Structural fault detection using a novelty measure. *Journal of Sound and Vibration* **1997**, *201*, 85–101. doi:10.1006/jsvi.1996.0747.

4. Corigliano, A.; Mariani, S. Parameter identification in explicit structural dynamics: Performance of the extended Kalman filter. *Computer Methods in Applied Mechanics and Engineering* **2004**, *193*, 3807–3835. doi:10.1016/j.cma.2004.02.003.

5. Eftekhar Azam, S.; Chatzi, E.; Papadimitriou, C. A dual Kalman filter approach for state estimation via output-only acceleration measurements. *Mechanical Systems and Signal Processing* **2015**, *60-61*, 866–886. doi:10.1016/j.ymssp.2015.02.001.

6. Farrar, C.; Worden, K. *Structural Health Monitoring A Machine Learning Perspective*; John Wiley & Sons, 2013.

7. Fink, O.; Wang, Q.; Svensen, M.; Dersin, P.; Lee, W.; Ducoffe, M. Potential, Challenges and Future Directions for Deep Learning in Prognostics and Health Management Applications. *Engineering Applications of Artificial Intelligence* **2020**, *92*, 103678. doi:10.1016/j.engappai.2020.103678.

8. Bishop, C.M. *Pattern Recognition and Machine Learning (Information Science and Statistics)*; Springer-Verlag, 2006.

9. LeCun, Y.; Bengio, Y.; Hinton, G. Deep learning. *Nature* **2015**, *521*, 436–444. doi:10.1038/nature14539.

10. Rosafalco, L.; Torzoni, M.; Manzoni, A.; Mariani, S.; Corigliano, A. Online structural health monitoring by model order reduction and deep learning algorithms. *Computers & Structures* **2021**, *255*, 106604. doi:10.1016/j.compstruc.2021.106604.

11. Sajedi, S.O.; Liang, X. Vibration-based semantic damage segmentation for large-scale structural health monitoring. *Computer-Aided Civil and Infrastructure Engineering* **2020**, *35*, 579–596. doi:10.1111/mice.12523.

12. Mirzazadeh, R.; Eftekhar Azam, S.; Mariani, S. Mechanical characterization of polysilicon MEMS: A hybrid TMCMC/POD-kriging approach. *Sensors (Switzerland)* **2018**, *18*, 1–19. doi:10.3390/s18041243.

13. Lam, H.F.; Yang, J.H.; Au, S.K. Markov chain Monte Carlo-based Bayesian method for structural model updating and damage detection. *Structural Control and Health Monitoring* **2018**, *25*, 1–22. doi:10.1002/stc.2140.

14. Torzoni, M.; Rosafalco, L.; Manzoni, A. A Combined Model-Order Reduction and Deep Learning Approach for Structural Health Monitoring Under Varying Operational and Environmental Conditions. Proceedings of the 7th International Electronic Conference on Sensors and Applications. MDPI: Basel, Switzerland, 2020, p. 94. doi:10.3390/ecsa-7-08258.

15. Meng, X.; Babaee, H.; Karniadakis, G.E. Multi-fidelity Bayesian Neural Networks: Algorithms and Applications. *Journal of Computational Physics* **2021**, *438*, 110361. doi:https://doi.org/10.1016/j.jcp.2021.110361.

16. Guo, M.; Manzoni, A.; Amendt, M.; Conti, P.; Hesthaven, J.S. Multi-fidelity regression using artificial neural networks: efficient approximation of parameter-dependent output quantities, 2021. https://arxiv.org/abs/2102.13403.

17. Meng, X.; Karniadakis, G.E. A composite neural network that learns from multi-fidelity data: Application to function approximation and inverse PDE problems. *Journal of Computational Physics* **2020**, *401*, 109020. doi:10.1016/j.jcp.2019.109020.

18. Quarteroni, A.; Manzoni, A.; Negri, F. *Reduced basis methods for partial differential equations: an introduction*; Springer, 2015.

19. Papadimitriou, C.; Lombaert, G. The effect of prediction error correlation on optimal sensor placement in structural dynamics. *Mechanical Systems and Signal Processing* **2012**, *28*, 105–127. doi:10.1016/j.ymssp.2011.05.019.

20. Hastings, W.K. Monte Carlo Sampling Methods Using Markov Chains and Their Applications. *Biometrika* **1970**, *57*, 97–109. doi:10.2307/2334940.

21. Goulet, J.A. *Probabilistic Machine Learning for Civil Engineers*; MIT Press, 2020.

22. Kingma, D.; Ba, J. Adam: A Method for Stochastic Optimization. International Conference on Learning Representations, 2015, Vol. 3, pp. 1–13. http://arxiv.org/abs/1412.6980.

23. Chollet, F.; others. Keras, 2015. https://keras.io.

24. Haario, H.; Saksman, E.; Tamminen, J. An adaptive Metropolis algorithm. *Bernoulli* **2001**, *7*, 223–242. doi:10.2307/3318737.