



Data driven Computational Mechanics at EXascale



Data driven Computational Mechanics at EXascale

Work program topic: EuroHPC-01-2019
Type of action: Research and Innovation Action (RIA)

Report on DMAP and CAE-based surrogate models

DELIVERABLE D2.3

Version No 1



<http://www.dcomex.eu/>

This project has received funding from the European High-Performance Computing Joint Undertaking Joint Undertaking ('the JU'), under Grant Agreement No 956201

DOCUMENT SUMMARY INFORMATION

Project Title	Data driven Computational Mechanics at EXascale
Project Acronym	DCoMEX
Project No:	956201
Call Identifier:	EuroHPC-01-2019
Project Start Date	01/04/2021
Related work package	WP 2
Related task(s)	Task 2.2
Lead Organisation	NTUA
Submission date	03/10/2023
Re-submission date	
Dissemination Level	PU

Quality Control:

	Who	Affiliation	Date
Checked by internal reviewer	George Stavroulakis	NTUA	01/10/2023
Checked by WP Leader	Vissarion Papadopoulos	NTUA	02/10/2023
Checked by Project Coordinator	Vissarion Papadopoulos	NTUA	02/10/2023

Document Change History:

Version	Date	Author (s)	Affiliation	Comment
1.0	26.09.2023	Ioannis Kalogeris	NTUA	



Table of Contents

Description	4
The diffusion maps algorithm	4
Theoretical background	4
Diffusion maps with variable-bandwidth kernels	5
Algorithmic implementation in the MSolve software	6
Surrogate modeling scheme using diffusion maps and neural networks	7
Numerical demonstration	8
Convolutional autoencoders	10
Theoretical background	10
Algorithmic implementation in the MSolve software	14
Surrogate modeling scheme using diffusion maps and neural networks	16
Numerical demonstration	17

1 Description

Deliverable 2.3 is associated to **WP2 “Surrogate modelling”** of the DCoMEX project, and it provides a comprehensive report on the development of two novel surrogate modeling scheme based (a) on the diffusion maps algorithm (DMAP) and (b) the machine learning algorithm of convolutional autoencoders (CAEs). This deliverable is an extension of Deliverable 2.2, where a set of numerical examples have been added to demonstrate the applicability of the proposed surrogate modeling techniques to complex engineering problems, as well as test their accuracy and assess the computational gains achieved. The outline of the deliverable is the following: Section 2 illustrates the theoretical background behind the DMAP algorithm, its algorithmic implementation in the MSolve software, a surrogate modeling strategy based on DMAP and a detailed numerical investigation. In a similar fashion, Section 3 presents the basic idea of CAEs, their algorithmic implementation in MSolve and the dedicated surrogate modeling strategy for complex engineering problems, which is then illustrated on a computational mechanics problem.

2 The diffusion maps algorithm

2.1 Theoretical background

Let $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_N]$ be a data set consisting of vectors $\mathbf{u}_i \in R^d$, which can be seen as N distinct realizations of an R^d -valued random variable and sampled independently with density $q(\mathbf{u})$. Next, assume a connectivity measure K between data pairs $\mathbf{u}_i, \mathbf{u}_j$ such as the Gaussian kernel

$$K_\epsilon(\mathbf{u}_i, \mathbf{u}_j) = \exp\left(\frac{-\left(\|\mathbf{u}_i - \mathbf{u}_j\|^2\right)}{4\epsilon}\right)$$

Next, a discrete approximation to the Laplacian L_ϵ is constructed as follows:

- Estimate the densities q_ϵ at the sample points \mathbf{u}_i as

$$q_\epsilon(\mathbf{u}_i) = \frac{1}{N} \sum_{j=1}^N K_\epsilon(\mathbf{u}_i, \mathbf{u}_j)$$

- Normalize the previously defined kernel K_ϵ as

$$\widehat{K}_\epsilon(\mathbf{u}_i, \mathbf{u}_j) = \frac{K_\epsilon(\mathbf{u}_i, \mathbf{u}_j)}{q_\epsilon(\mathbf{u}_i)^\alpha q_\epsilon(\mathbf{u}_j)^\alpha}$$

Where for $\alpha = 1$ the discrete Laplacian approximates the Laplace-Beltrami operator, while $\alpha = 1/2$ approximates a diffusion operator.

- Estimate the new densities

$$\widehat{q}_\epsilon(\mathbf{u}_i) = \frac{1}{N} \sum_{j=1}^N \widehat{K}_\epsilon(\mathbf{u}_i, \mathbf{u}_j)$$

- If we define the matrix $\mathbf{K} = [K_{ij}] = \widehat{K}_\epsilon(\mathbf{u}_i, \mathbf{u}_j)$ and the diagonal matrix $\mathbf{D} = [D_{ii}] = q_\epsilon(\mathbf{u}_i)$, then the discrete approximation of the weighted Laplacian is given by the expression:

$$\mathbf{L}_\varepsilon = \frac{\mathbf{D}^{-1}\mathbf{K} - \mathbf{I}_N}{\varepsilon}$$

The solution to the eigenvalue problem $\mathbf{L}_\varepsilon \boldsymbol{\Psi} = \lambda \boldsymbol{\Psi}$ will produce the sequence of eigenvalues $0 = \lambda_0 \geq \lambda_1 \geq \lambda_2 \geq \dots$ and right eigenvectors $\boldsymbol{\Psi}_j$ for the operator. In practice, only the first n non-trivial eigenvectors are kept with n obtained from the expression

$$n = \operatorname{argmin}_{n, n \geq 2} \left(\frac{\lambda_1}{\lambda_n} < \operatorname{tol} \right)$$

Then, the diffusion map operator $\Psi_\varepsilon: u \rightarrow R^n$ can be defined as

$$\Psi_\varepsilon(\mathbf{u}) = [e^{\lambda_1 \varepsilon} \psi_1(\mathbf{u}), e^{\lambda_2 \varepsilon} \psi_2(\mathbf{u}), \dots, e^{\lambda_n \varepsilon} \psi_n(\mathbf{u})]$$

2.2 Diffusion maps with variable-bandwidth kernels

In several data-driven applications, the samples follow some distribution which is unknown a priori. It is expected that the samples belonging to the tails of the distribution will be fewer and, thus, there will be regions on the manifold that will be more sparsely delineated. To address this issue in classical kernel methods the idea of the variable-bandwidth (or self-tuning) kernels has been proposed and illustrated herein. The main differentiation with respect to the classical DMAP algorithm lies in the form of the kernel used, which in this setting becomes:

$$K_\varepsilon^{VB}(\mathbf{u}_i, \mathbf{u}_j) = \exp\left(\frac{-\left(\|\mathbf{u}_i - \mathbf{u}_j\|^2\right)}{4\varepsilon\rho(\mathbf{u}_i)\rho(\mathbf{u}_j)}\right)$$

Following the construction for the graph Laplacian of the previous sections, in this case the sample densities are

$$q_\varepsilon^{VB}(\mathbf{u}_i) = \sum_{j=1}^N \frac{K_\varepsilon(\mathbf{u}_i, \mathbf{u}_j)}{\rho(\mathbf{u}_i)^m}$$

which are used to construct the kernel

$$K_{\varepsilon, \alpha}^{VB}(\mathbf{u}_i, \mathbf{u}_j) = \frac{K_\varepsilon^{VB}(\mathbf{u}_i, \mathbf{u}_j)}{q_\varepsilon^{VB}(\mathbf{u}_i)^\alpha q_\varepsilon^{VB}(\mathbf{u}_j)^\alpha}$$

Setting $q_{\varepsilon, \alpha}^{VB}(\mathbf{u}_i) = \sum_{j=1}^N K_{\varepsilon, \alpha}^{VB}(\mathbf{u}_i, \mathbf{u}_j)$, we can obtain the normalized kernel

$$\widehat{K}_{\varepsilon, \alpha}^{VB}(\mathbf{u}_i, \mathbf{u}_j) = \frac{K_{\varepsilon, \alpha}^{VB}(\mathbf{u}_i, \mathbf{u}_j)}{q_{\varepsilon, \alpha}^{VB}(\mathbf{u}_i)}$$

and the weighted Laplacian for this formulation becomes

$$L_{\varepsilon, \alpha}^{VB}(\mathbf{u}_i, \mathbf{u}_j) = \frac{\widehat{K}_{\varepsilon, \alpha}^{VB}(\mathbf{u}_i, \mathbf{u}_j) - \delta_{ij}}{\varepsilon \rho(\mathbf{u}_i)^2}$$

2.3 Algorithmic implementation in the MSolve software

The code for implementing the variable-bandwidth diffusion maps algorithm can be found in <https://github.com/mgroupntua/MSolve.MachineLearning>¹. In particular, the C# class `DiffusionMapsAlgorithm.cs` in the `MGroup.MachineLearning` folder implements the aforementioned procedure for an input data set. An example illustrating the use of this class is provided in the `MGroup.MachineLearning.Tests` folder, called `DMAPExample.cs`.

In this particular example, an initial data set is considered which consists of 2000 points in R^2 , generated from a 2-dimensional Gaussian distribution centered at zero with covariance $C = 0.04I_2$. Using the syntax outlined below, a new object called `DMAPE` from the `DiffusionMapsAlgorithm` class is generated, taking as input from the user a specified set of variables. Then the method `ProcessData()` applies the DMAPE algorithm and computes the member variables `DMAPE.DMAPEigenvalues[.]` and `DMAPE.DMAPEigenvalues[.]`.

- `dataSet` : the initial data set
- `numberOfKNN`: number of k-nearest neighbors used in the evaluation of the kernel $K_{\varepsilon, \alpha}^{VB}(\mathbf{u}_i, \mathbf{u}_j)$
- `numberOfKDE`: number of k-nearest neighbors required to estimate the kernel parameter ε
- `differentialOperator`: 1 – Laplace Beltrami operator, 2- generator of grad systems
- `numberOfEigenvectors` : The number of eigenvectors requested by the user

```
DiffusionMapsAlgorithm DMAPE = new DiffusionMapsAlgorithm(dataSet,  
numberOfKNN, NNofKDE, differentialOperator, numberOfEigenvectors);  
  
DMAPE.ProcessData();
```

The data used in this particular example are shown in figure 1, while figure 2 depicts the first 10 non-trivial DMAPE eigenvalues.




Figure 1: initial data samples

¹ The code has originally been submitted in the repo <https://github.com/YiannisKalogeris/MSolve.MachineLearning>




Figure 2: The first 10 diffusion map eigenvalues

2.4 Surrogate modeling scheme using diffusion maps and neural networks

Consider the modeling of a parametrized physical system governed by partial differential equations:

$$\begin{aligned} \frac{\partial u(\mathbf{x}, t; \boldsymbol{\theta})}{\partial t} + \mathcal{L}[u(\mathbf{x}, t; \boldsymbol{\theta})] &= f(\mathbf{x}, t; \boldsymbol{\theta}), & \mathbf{x} \in \Omega, t \in [0, T], \boldsymbol{\theta} \in \Theta \\ \mathcal{B}[u(\mathbf{x}, t; \boldsymbol{\theta})] &= b(\mathbf{x}, t; \boldsymbol{\theta}), & \mathbf{x} \in \partial \Omega, t \in [0, T], \boldsymbol{\theta} \in \Theta \\ \mathcal{C}[u(\mathbf{x}, 0; \boldsymbol{\theta})] &= c(\mathbf{x}; \boldsymbol{\theta}), & \mathbf{x} \in \Omega, \boldsymbol{\theta} \in \Theta \end{aligned}$$

where $u(\mathbf{x}, t; \boldsymbol{\theta})$ is the (scalar) field of interest, \mathcal{L} is a general differential operator that involves spatial derivatives and $f(\mathbf{x}, t; \boldsymbol{\theta})$ is a source field. Also, \mathcal{B} is the operator for the boundary conditions defined on $\partial \Omega$, \mathcal{C} is the operator for the initial conditions of the problem at $t = 0$ and $\boldsymbol{\theta} \in \Theta$ is the vector of the problem's parameters.

The discrete solution to the above set of equations for a given parameter instance $\boldsymbol{\theta}$ is obtained using numerical discretization techniques, such as the finite element method. In the context of parametric simulations, a set of parameter samples $\{\boldsymbol{\theta}_j\}_j^N$ is initially generated and the discretized solution vectors $\mathbf{u}_j(t_k) \in \mathbb{R}^d$ corresponding to each parameter instance $\boldsymbol{\theta}_j$ and time instance t_k are obtained through the numerical solution of the partial differential equation.

Based on the above, the numerical solution of the PDE can be regarded as a nonlinear mapping from the parametric space of \mathbb{R}^d . However, for a large-scale ($d \gg 1$) nonlinear dynamic problem, solving the corresponding numerical model can take significant computational resources and memory requirements. An intuitive approach to reduce this cost would be to exchange this mapping with another one that will be adequately accurate and much faster to compute. For this task, a feed-forward neural network, or FFNN for short trained over a reduced set of system solution would seem an appealing choice. Nevertheless, a FFNN with such high-dimensional output translates to a massive number of network parameters to be calibrated and the associated computational cost and memory demands would render this process unprofitable.

Our approach to overcome this issue is by employing the DMAP algorithm in order to reduce the dimensionality of the model's output, which will allow us to efficiently train a FFNN that will map points from the parametric space to their reduced representation in the DMAP space. Next, another mapping needs to be established from the DMAP space to the high-dimensional solution space that will reconstruct the system's solutions from their reduced representations. For this purpose, in our investigation we chose the Laplacian Pyramids interpolation algorithm. The steps for the algorithmic implementation of the proposed procedure are outlined below.

1. Generate N samples of the vector of random parameters $\boldsymbol{\theta}$ and perform N simulations. Collect the solutions at N_t time increments to obtain the $d \times N_{tr}$ matrix of solutions (snapshots) $\mathbf{U} = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_{N_{tr}}]$, where $N_{tr} = N \times N_t$, N_t being the total number of time steps in the problem.

2. Given the matrix \mathbf{U} of the N_{tr} distinct solutions, utilize the machinery provided by the diffusion maps algorithm to identify a low dimensional representation $\mathbf{u} \in \mathbb{R}^d \rightarrow \mathbf{z} \in \mathbb{R}^n$, with $n \ll d$.
3. Train a FFNN using $\{\mathbf{e}_i\}_{i=1}^{N_{tr}} = \{(\boldsymbol{\theta}_j, t_k)\}_{1 \leq j \leq N, 0 \leq k \leq N_t}$ as input and the DMAP coordinates $\{\mathbf{z}_i\}_{i=1}^{N_{tr}}$ as output.
4. For a new vector of parameter values $\mathbf{e} = (\boldsymbol{\theta}, t)$ utilize the previously trained FFNN nonlinear mapping to obtain its image in the DMAP space as $FFNN(\mathbf{e}) = \mathbf{z} \in \mathbb{R}^n$.
5. Solve the pre-image problem with the Laplacian pyramid interpolation scheme to obtain the solution $\mathbf{u} = LP(\mathbf{z}) \in \mathbb{R}^d$

In the above, steps 1-3 constitute the offline (training) phase of the surrogate, while steps 4-5 are the online phase and these can be repeated for as many parameter instances as required for the purposes of the analysis. A schematic representation of the proposed surrogate modeling strategy is given in figure 3.




Fig 3. Concept of the proposed surrogate modeling strategy

2.5 Numerical Demonstration

In this section, we consider the 2 – D convection-diffusion-reaction equation given below, as a test case for the DMAP-based surrogate modeling strategy.

$$\frac{\partial u}{\partial t} + \boldsymbol{\alpha} \cdot \nabla u - \nabla \cdot (\nu \nabla u) + ru = 0 \quad \text{in } \Omega \times [0, T]$$

$$u(\mathbf{x}, t) = 0 \quad \text{on } \partial\Omega$$

$$u(\mathbf{x}, 0) = u_0 \quad \text{on } \Omega$$

where, $\Omega = [-0.5, 0.5] \times [-0.5, 0.5]$ and $T = 2\pi$. The boundary condition u_0 for this problem is:

$$u_0(x, y) = \begin{cases} \frac{1}{4} \left(1 + \cos\left(5\pi x - \frac{5\pi}{6}\right)\right) \left(1 + \cos\left(5\pi y - \frac{5\pi}{6}\right)\right), & \left(x - \frac{1}{6}\right)^2 + \left(y - \frac{1}{6}\right)^2 < 0.2^2 \\ 0, & \text{otherwise} \end{cases}$$

The velocity field is given by $\mathbf{a} = (c_1 y, c_2 x)$ with c_1, c_2 being random parameters following the Gaussian distribution $N(0, 0.2)$. In addition, the diffusion coefficient ν and the reaction coefficient r follow lognormal distributions with mean values 0.001 and 0.1, respectively. The coefficient of variation is taken equal to 0.3 for both random variables.

For the numerical solution of the above equation, the finite element method was employed for the spatial discretization using a 40×40 bilinear mesh, leading to 1681 unknowns (degrees of freedom), while for the time discretization the Crank-Nicolson scheme was used with $dt = 6\pi/100$. In order, to apply the proposed methodology, a set of 300 system solutions were first obtained, which were stored in a matrix to construct the initial data set. By applying the DMAP algorithm on the data set, the value $\varepsilon = 0.0583$ for the kernel parameter was selected, based on the plot of the eigenvalues of figure 4. This value led to 10 eigenvectors for the low-dimensional representation of the data set.




Fig 4. Selection of ε and number of eigenfunctions for the DMAP representation

Next, we used the 30000 labeled data pairs $\{\mathbf{e}_i\}_{i=1}^{30000} = \{(v_j, r_j, c_{1,j}, c_{2,j}, t_k)\}_{1 \leq j \leq 300, 1 \leq k \leq 100}$ to train a FFNN using these as input and the corresponding 10-dimensional DMAP coordinates as output. In addition, we implemented the Laplacian pyramid interpolation using a 2-stage scheme to map the DMAP coordinates to the high-dimensional solution space. With these tools, the training process of the surrogate was completed and, subsequently, we used it to perform a Monte Carlo analysis with 2000 simulations. In the following figures the results obtained by the surrogate are compared to those of the high-fidelity model to verify its accuracy.

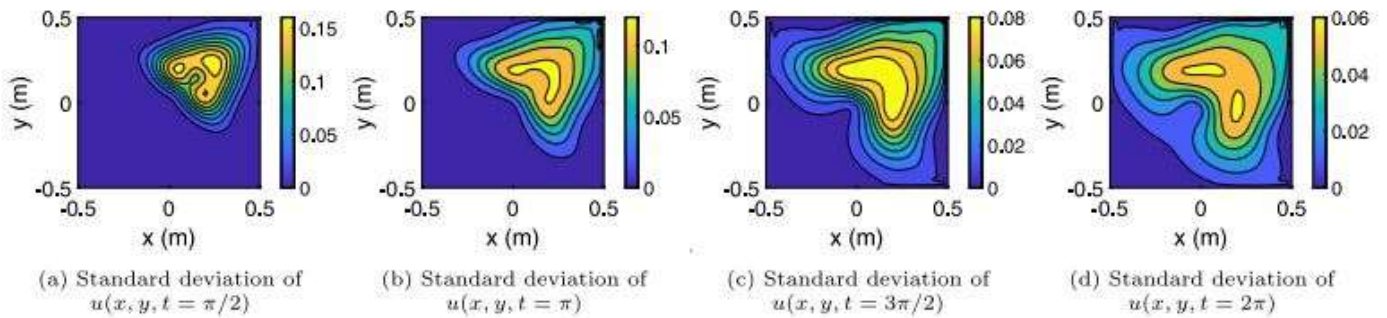


Fig 5: Standard deviation of $u(x, y, t)$ at specific instants obtained from the exact model.



DCoMEX

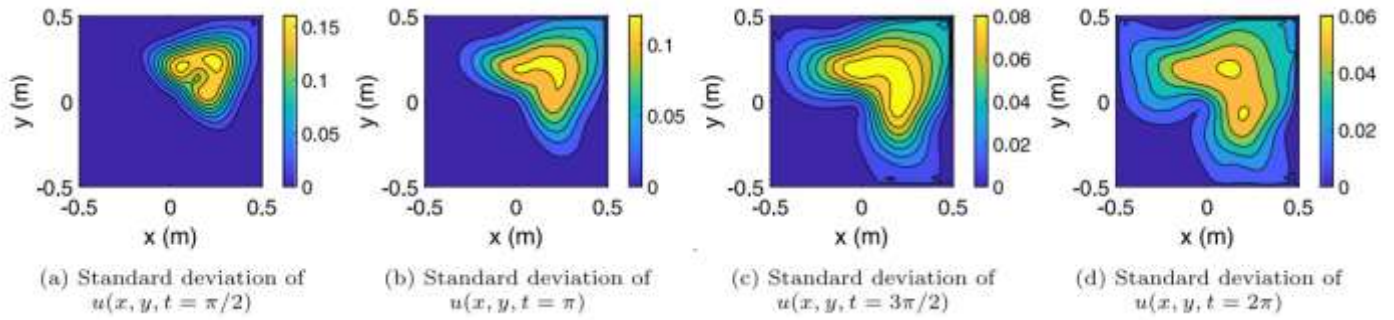


Fig 6: Standard deviation of $u(x, y, t)$ at specific instants obtained from the surrogate model.

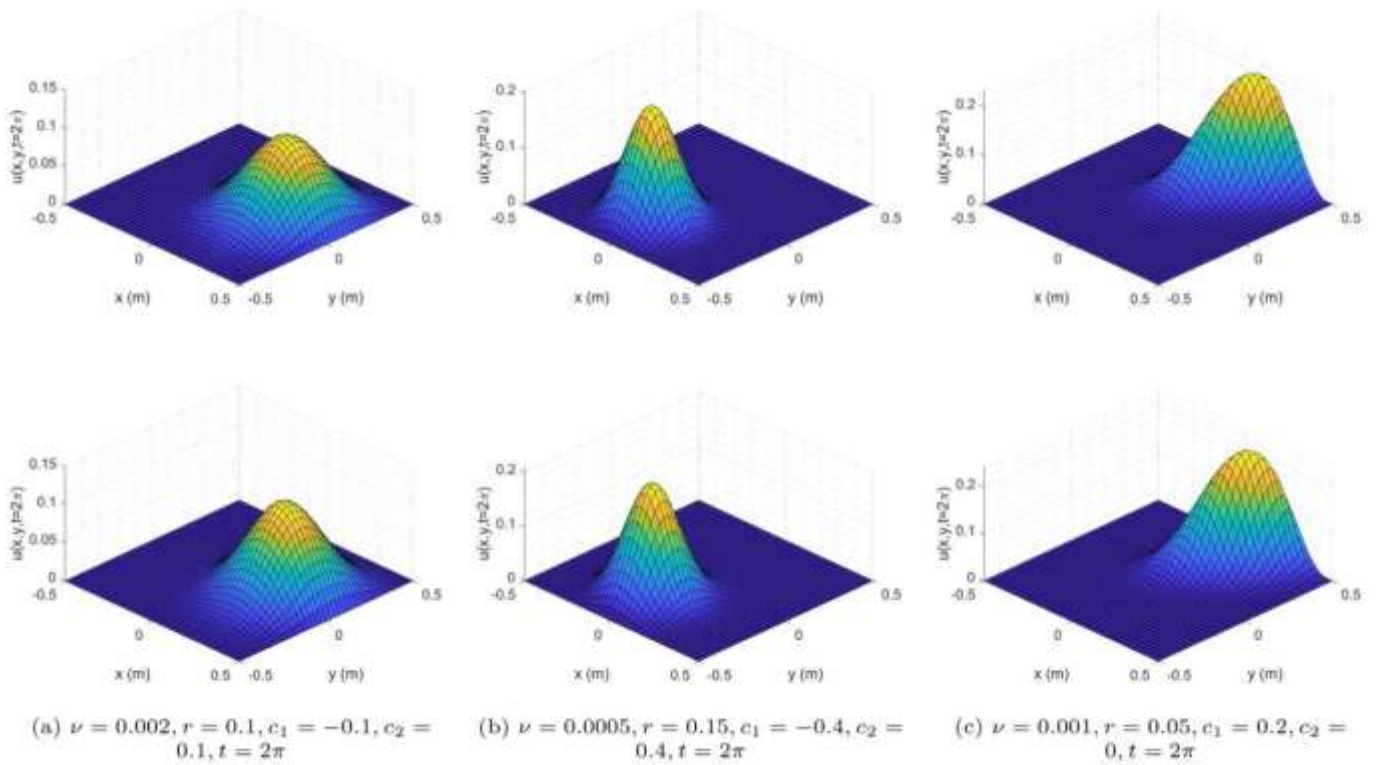


Fig 7: Exact solutions (first row) and surrogate's predictions (second row) for specified parameter values.

3 Convolutional autoencoders

3.1 Theoretical background

The Autoencoder (AE) concept is regarded as a neural network that learns from an unlabeled data set in an unsupervised manner. The aim of an AE is to learn a reduced representation for a set of data, known as encoding, and then learn how to reconstruct the original input from the encoded input with the minimum possible error. The latter part of the AE is called the decoder.

In particular, let \mathbf{X} be a subset of \mathbb{R}^d with $\mathbf{x} \in \mathbf{X}$ denoting an element of the set. Then, the AE's encoder and decoder are defined as transition maps φ, ψ such that:

$$\begin{aligned} \varphi: \mathbf{X} \subseteq \mathbb{R}^d &\rightarrow \mathbf{H} \subseteq \mathbb{R}^l \\ \psi: \mathbf{H} \subseteq \mathbb{R}^l &\rightarrow \mathbf{X} \subseteq \mathbb{R}^d \\ \varphi, \psi &= \operatorname{argmin}_{\varphi, \psi} \|\mathbf{X} - (\psi \circ \varphi)\mathbf{X}\| \end{aligned}$$

with the dimension l typically being much smaller than d .

Now, let us consider the simplest case, where the encoder has only one hidden layer. It takes an input $\mathbf{x} \in \mathbb{R}^d$ and sends it to $\mathbf{h} = \varphi(\mathbf{x}) \in \mathbb{R}^l$ with $\mathbf{h} = \sigma(\mathbf{W}\mathbf{x} + \mathbf{b})$, σ being an activation function (e.g. tanh or ReLU), \mathbf{W} a weight matrix and \mathbf{b} a bias vector. The image \mathbf{h} of \mathbf{x} is the latent or encoded representation of \mathbf{x} and \mathbf{H} is the latent or feature space.

The decoder's task is to establish the reverse mapping ψ that will reconstruct the input \mathbf{x} , given its latent representation \mathbf{h} . Again, considering a one-hidden layer, the reconstructed point $\tilde{\mathbf{x}} = \psi(\mathbf{h})$ is given by: $\tilde{\mathbf{x}} = \tilde{\sigma}(\tilde{\mathbf{W}}\mathbf{h} + \tilde{\mathbf{b}})$, with $\tilde{\sigma}, \tilde{\mathbf{W}}, \tilde{\mathbf{b}}$ being different than those of the encoder. Also, the network's architecture selected for the encoder can be different than that of the decoder and the number of hidden layers can be greater than one, leading to the so-called deep AEs.

AEs are trained by a back propagation algorithm, which is the most commonly used algorithm for the training of NNs. The training loss function becomes the reconstruction error between the input points \mathbf{x}_i and their respective output $\tilde{\mathbf{x}}_i$, that is:

$$L = \frac{1}{N} \sum_{i=1}^N \|\mathbf{x}_i - \tilde{\mathbf{x}}_i\|$$

Despite their powerful dimensionality reduction properties, AEs face significant challenges when dealing with very high-dimensional inputs, due to the fact that the number of trainable parameters increases drastically with an increase in the input's dimensionality. In addition, AEs are not capable of capturing the spatial features of the input (e.g. when dealing with images) nor the sequential information in the input (e.g. when dealing with sequence data). To remedy these issues, a new type of AEs has emerged, that of convolutional autoencoders (CAEs). Similarly to AEs, CAEs also consist of an encoder and a decoder that are trained to minimize the loss function, but they are built from different layer types. Specifically, in CAEs the encoder part is built using a combination of convolutional layers, fully connected layers, pooling layers and normalization layers, while the decoder is built from deconvolutional layers and unpooling layers along with fully connected and normalization layers. Intuitively, CAEs can be viewed as extensions of ordinary AEs in the same way that CNNs are extensions of FFNNs.

Convolutional layers take as input a $n - D$ array \mathbf{M} and apply a filter (a.k.a kernel) F of specified size to the elements of \mathbf{M} in a moving window fashion. This process is schematically depicted in following figure.




Fig. 8: Schematic representation of a 2 – D convolutional filter with strides $s_h = 2$ and $s_v = 2$

To better clarify this process, let us consider a 2 – D array $\mathbf{M} = [m_{ij}]$ and its encoded version $\mathbf{M}^{enc} = [\mu_{ij}]$, called feature map, which is obtained after applying a filter $\mathbf{W} = [w_{ij}]$ of size $f_h \times f_w$, moving with vertical stride s_v and horizontal stride s_h . The element μ_{ij} of \mathbf{M}^{enc} is given by the equation:

$$\mu_{ij} = \sum_{u=1}^{f_h} \sum_{v=1}^{f_w} m_{i'j'} \cdot w_{uv} + b \text{ with } i' = i \times s_v + u, j' = j \times s_h + v$$

where b is the bias term and w_{uv} is the element of the filter \mathbf{W} that gives the connection weight between elements of \mathbf{M}^{enc} and the element of \mathbf{M} within the respective window. This layer architecture is significantly more economical than that of a fully connected layer since the parameters involved are only the $f_h \times f_w$ elements of the filter \mathbf{W} and the bias term b . The filter parameters do not require to be manually defined, instead the convolutional layer will learn the most appropriate filter for the task. Also, a convolutional layer can have multiple filters, in which case it outputs one feature map \mathbf{M}_k^{enc} per each filter k . We shall write $\mathbf{M}^{enc} = ConvNN(\mathbf{M})$ to denote the application of several convolutional layers, with multiple filters each, to an array \mathbf{M} .

On the other hand, a deconvolutional layer performs the reverse operation of convolution, called deconvolution, and it is used to construct decoding layers. Their function is to multiply each input value by a filter elementwise. For instance a 2D $f_h \times f_w$ deconvolution filter maps a 1×1 spatial region of the input to an $f_h \times f_w$ region of the output. Thus, the filters learned in the deconvolutional layers create a basis used for the reconstruction of the inputs' shape, taking into consideration the required shape of the output. As before, a deconvolutional layer can have multiple filters, while several deconvolutional layers can be stacked for building deep architectures. The decoding procedure can be represented as: $\mathbf{M} = DeconvNN(\mathbf{M}^{enc})$.

Based on the above, the CAE's architecture consists of convolutional, deconvolutional and dense layers and is typically used for dimensionality reduction and reconstruction purposes. In practice, the CAE's encoder uses a number of convolutional layers to compress the input and once the desirable level of reduction has been achieved, the encoded matrix is flattened into a vector. Then, a dense layer is employed to map this vector to its latent representation. In the reverse direction, the decoder starts by taking the latent representation and transforming it into a vector through a denser layer. Subsequently, the input reconstruction is achieved by the deconvolutional layers. Thus, the loss for CAEs becomes:

$$L = \frac{1}{N} \sum_{i=1}^N \|\mathbf{M}_i - \widetilde{\mathbf{M}}_i\|$$

where \mathbf{M}_i denotes the input arrays used for training and $\widetilde{\mathbf{M}}_i = DeconvNN(ConvNN(\mathbf{M}))$ the corresponding CAE's output. A schematic representation of a deep CAE is presented in figure 5.

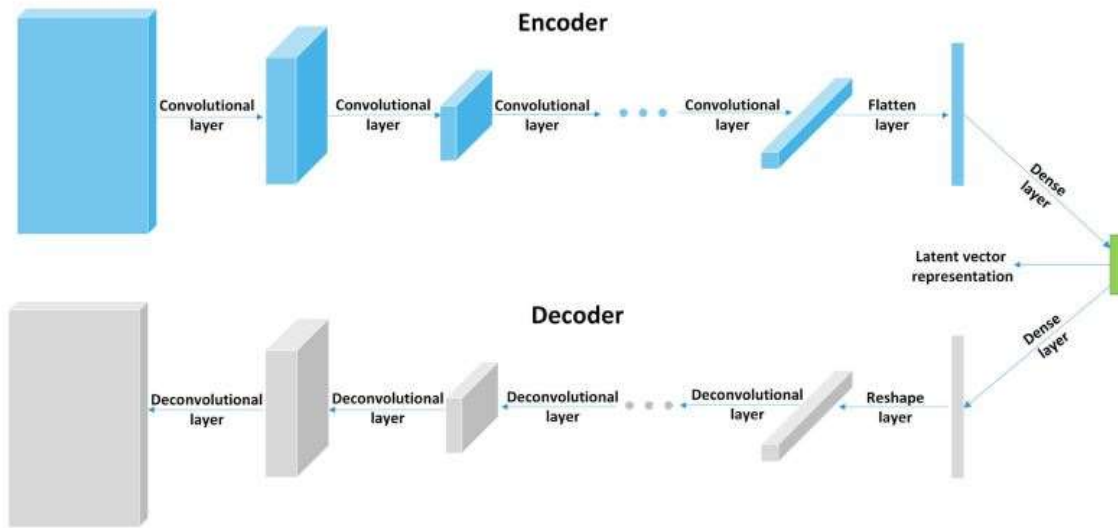


Fig.9: Schematic representation of a deep convolutional autoencoder

Lastly, aside of convolutional, deconvolutional and dense layers, two other important layer types often employed in CAEs are those of pooling and unpooling. Pooling layers are quite similar to convolutional layers in the sense that they downsample the input in order to decrease its size, however, they do not involve any trainable parameters. Their goal is to reduce the computational load, the memory usage, and the number of parameters. Common types of pooling layers include the max pooling layer and the average pooling layer. The first outputs the maximum value from the portion of the input covered by the filter and all other inputs are neglected. Accordingly, average pooling layers return the average from the portion of the input. On the other hand, unpooling layers perform the reverse operation of pooling and their aim is to reconstruct the original size of each rectangular patch. These operations are schematically depicted in figure 6.

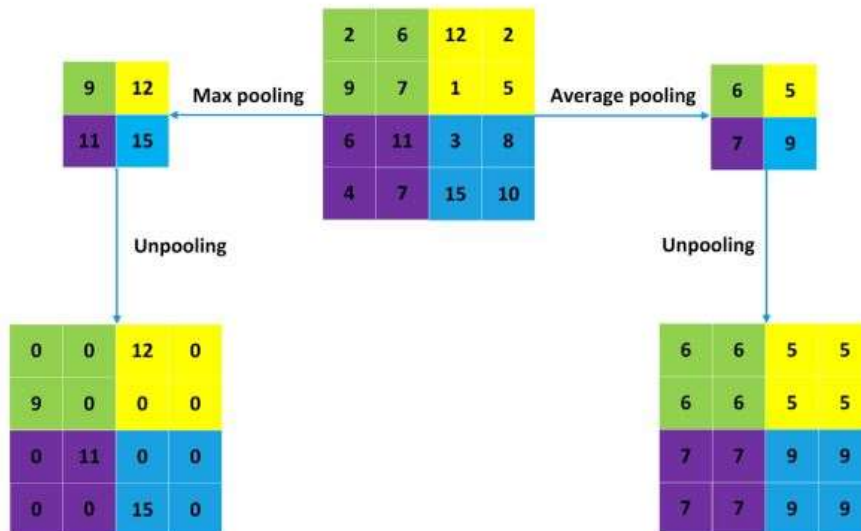


Fig. 10: Examples of pooling and unpooling



3.2 Algorithmic implementation in the MSolve software

The code regarding the Convolutional Autoencoders is located in <https://github.com/mgroupntua/MSolve.MachineLearning>. Specifically, the C# class **ConvolutionalAutoencoder** in the folder **NeuralNetworks** of the project **MGroup.MachineLearning.Tensorflow** can be readily used to implement a CAE.

ConvolutionalAutoencoder Constructor

ConvolutionalAutoencoder(normalization, optimizer, lossFunc, encoderLayers, decoderLayers, epochs, batchSize, seed, classification)

normalization: normalization option for the preprocessing stage of the CAE (e.g. MinMaxNormalization, ZScoreNormalization)

optimizer: optimization algorithm used for the CAE training (e.g. Adam, SGD, RMSProp)

lossFunc: loss function which will be minimized in the optimization procedure (e.g. MeanSquaredError, MeanAbsoluteError)

encoderLayers: network architecture regarding the layers of the encoder (e.g. Convolutional2DLayer, MaxPooling2DLayer)

decoderLayers: network architecture regarding the layers of the decoder (e.g. ConvolutionalTranspose2DLayer, UpSampling2DLayer)

epochs: total number of optimization iterations in terms of processing the full dataset

batchSize: total number of samples that will be processed in a single backpropagation step (default: 32)

seed: (*optional*)seed of the random number generator (default: null)

classification: controls the nature of the prediction i.e. for classification or regression (default: false)

ConvolutionalAutoencoder Methods

Train(trainX, testX)

trainX: train dataset (*remark: for an AE the input and output are identical*)

testX: (*optional*)test dataset (default: null)

method objective: to train a CAE with the provided train dataset

ValidateNetwork(testX)

testX: test dataset

method objective: to validate a trained CAE with the provided test dataset



`MapFullToReduced(initialStimuli)`

initialStimuli: collection of samples in the initial/full representation

method objective: to calculate the representation of the provided samples in the reduced space

`MapReducedToFull(reducedStimuli)`

reducedStimuli: collection of samples in the reduced/latent representation

method objective: to calculate the representation of the provided samples in the initial space

`EvaluateResponses(initialStimuli)`

reducedStimuli: collection of samples in the initial/full representation

method objective: to calculate the representation of the provided samples in the initial space after performing the full CAE procedure i.e. encoding and decoding

`SaveNetwork(netPath, weightsPath, normalizationPath)`

netPath: path of the operating system where the CAE architecture is saved to

weightsPath: path of the operating system where the CAE trained parameters are saved to

normalizationPath: path of the operating system where the normalization information is saved

method objective: to save the necessary information of after the training of a CAE, so it can be used in the future or shared

`LoadNetwork(netPath, weightsPath, normalizationPath)`

netPath: path of the operating system where the CAE architecture is loaded from

weightsPath: path of the operating system where the CAE trained parameters are loaded from

normalizationPath: path of the operating system where the normalization information is loaded from

method objective: to load the necessary information of an already trained CAE

ConvolutionalAutoencoder Example

```
class ConvolutionalAutoencoderTest
```

class objective: to illustrate the training accuracy of a CAE in the case of image classification. Specifically, it verifies the correct implementation and the accuracy of the CAE that has been trained on a subset of the MNIST dataset. Both the reduced representation and the reconstruction of the images are monitored.

3.3 Surrogate modeling scheme using convolutional autoencoders and neural networks

In a similar fashion to section 2.4, let us consider again the modeling of a parametrized physical system governed by partial differential equations:

$$\begin{aligned} \frac{\partial u(\mathbf{x}, t; \boldsymbol{\theta})}{\partial t} + \mathcal{L}[u(\mathbf{x}, t; \boldsymbol{\theta})] &= f(\mathbf{x}, t; \boldsymbol{\theta}), & \mathbf{x} \in \Omega, t \in [0, T], \boldsymbol{\theta} \in \Theta \\ \mathcal{B}[u(\mathbf{x}, t; \boldsymbol{\theta})] &= b(\mathbf{x}, t; \boldsymbol{\theta}), & \mathbf{x} \in \partial \Omega, t \in [0, T], \boldsymbol{\theta} \in \Theta \\ \mathcal{C}[u(\mathbf{x}, 0; \boldsymbol{\theta})] &= c(\mathbf{x}; \boldsymbol{\theta}), & \mathbf{x} \in \Omega, \boldsymbol{\theta} \in \Theta \end{aligned}$$

By applying spatial and temporal discretization techniques, the complete time-history response of the system is given by $\mathbf{U}(\boldsymbol{\theta}) = [\mathbf{u}(\boldsymbol{\theta}, t_1), \mathbf{u}(\boldsymbol{\theta}, t_2), \dots, \mathbf{u}(\boldsymbol{\theta}, t_{N_t})] \in \mathbb{R}^{d \times N_t}$, where N_t is the number of time increments in the temporal discretization.

The surrogate modeling approach proposed in this section is based on the powerful dimensionality reduction capabilities of CAEs. To this purpose, the PDEs are solved with the classic numerical procedure for a small, yet sufficient number, N , of parameter values in order to obtain a data set of time history matrices $\{\mathbf{U}_i\}_{i=1}^N$. The CAE (encoder and decoder) is trained over this data set minimizing the reconstruction mean square error. The encoded representation of each time history solution matrix \mathbf{U}_i is a low dimensional vector $\mathbf{z}_i \in \mathbb{R}^l$ ($l \ll d$), which allows a FFNN to be trained accurately and efficiently in order to construct a mapping between the PDEs parametric space and the encoded solution space. It should be mentioned that the optimal architecture and hyperparameters of the CAE and FFNN are typically obtained via a trial and error procedure. After the training phase is completed, the proposed surrogate scheme works as follows. For a new input parameter vector, the encoded vector representation of the time history solution matrix is calculated by the FFNN and, subsequently, the entire time history matrix is delivered by the CAE's decoder. This way a large number of additional simulations can be performed at minimum computational cost.

The steps for the algorithmic implementation of the proposed procedure are outlined below.

1. Generate N samples of the vector of random parameters $\boldsymbol{\theta}$ and perform N simulations. Collect the solutions in three-dimensional array $N \times d \times N_t$, where d is the number of degrees of freedom and N_t the number of time increments.
2. Train a CAE over the N time history solution matrices $\mathbf{U}_i \in \mathbb{R}^{d \times N_t}$, collected at step 1, to obtain the encoded low dimensional vector representations $\mathbf{z}_i \in \mathbb{R}^l$ of these matrices along with the reconstruction map.
3. Train a FFNN to establish a mapping from the parameter space to the low dimensional encoded space, that is $\mathbf{z}_i = \text{FFNN}(\boldsymbol{\theta}_i)$.
4. For a new vector of parameter values $\boldsymbol{\theta}$ utilize the trained FFNN to obtain the encoder vector representation $\mathbf{z} \in \mathbb{R}^l$ of the solution matrix.
5. The CAE's decoder is used to produce the solution matrix \mathbf{U} based the encoded representation \mathbf{z} of the previous step.

In the above, steps 1-3 constitute the offline (training) phase of the surrogate, while steps 4-5 are the online phase and these can be repeated for as many parameter instances as required for the purposes of the analysis. A schematic representation of the proposed surrogate modeling strategy is given in figures 11 and 12.

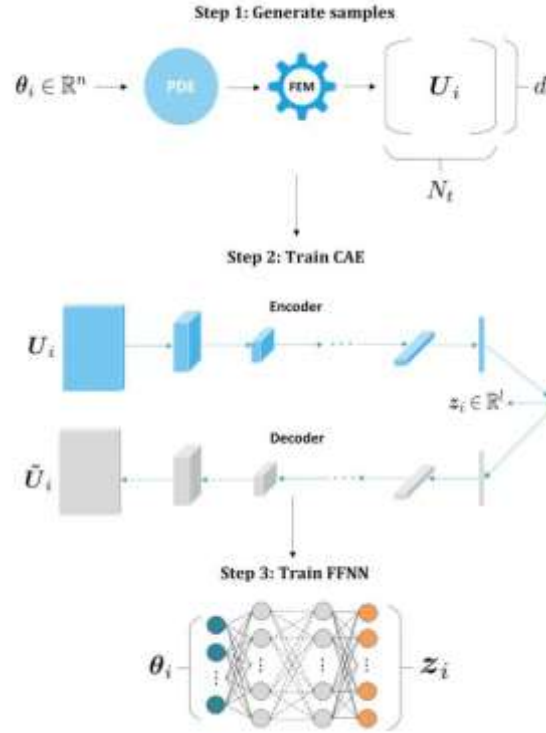


Fig.11: Offline phase of the proposed surrogate modeling method

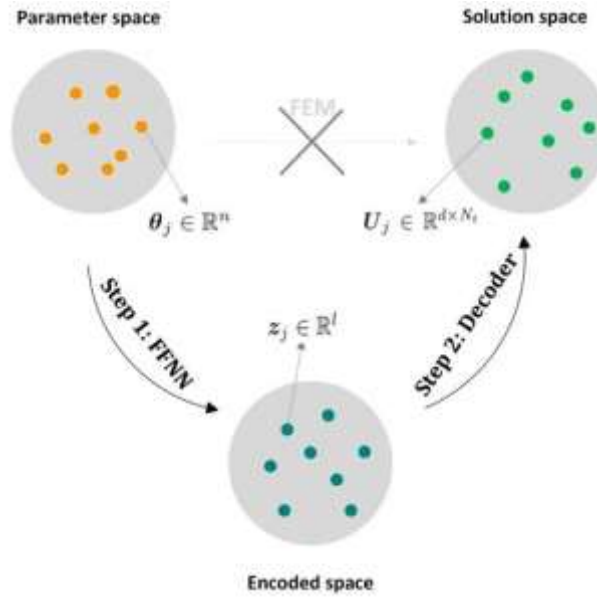


Fig. 12: Online phase of the proposed surrogate modeling method

3.4 Numerical Demonstration

In this section, we consider the 2 – D elasticity problem of coupled shear walls under seismic loading given below, as a test case for the CAE-based surrogate modeling strategy.

$$\begin{aligned} \frac{E}{2(1+\nu)} \nabla^2 u_x + \frac{E}{2(1-\nu)} \frac{\partial}{\partial x} \left(\frac{\partial u_x}{\partial x} + \frac{\partial u_y}{\partial y} \right) &= \rho \frac{\partial^2 u_x}{\partial t^2} \\ \frac{E}{2(1+\nu)} \nabla^2 u_y + \frac{E}{2(1-\nu)} \frac{\partial}{\partial y} \left(\frac{\partial u_x}{\partial x} + \frac{\partial u_y}{\partial y} \right) &= \rho \frac{\partial^2 u_y}{\partial t^2} \end{aligned}$$

where, $u_x = u_x(x, y, t)$, $u_y = u_y(x, y, t)$ are the displacement fields, E is the modulus of elasticity, ν is the Poisson ratio and ρ is the material's density. The geometry and the finite element mesh of the problem are shown in figure 13.




Fig. 13 Geometry and finite element mesh of the coupled shear walls problem

For this problem, we assumed that each story has each own modulus of elasticity, i.e. E_1, E_2, E_3 , which are all considered uncorrelated random variables, following a lognormal distribution with mean value 30GPa and standard deviation 7.5GPa. The structure is subjected to a seismic loading based on the accelerogram of a previous earthquake (1972 Kefalonia earthquake) with a duration of 6.00 seconds.

To apply the proposed methodology, initially, a set of 500 triplets of parameters $\{[E_1^i, E_2^i, E_3^i]\}_{i=1}^{500}$ are generated and the structure is solved using the finite element method using a spatial mesh of 876 quadrilateral elements for the spatial discretization leading to a total of 1966 degrees of freedom and 600 time steps for the time discretization. Thus, the solution matrices $\{U_i\}_{i=1}^{500} \in R^{1966 \times 600}$ are obtained, corresponding to each of the 500 realizations of the parameter triplets. These matrices of 'exact' solutions served as the data set over which the CAE is trained to find a low-dimensional representation of only $l = 64$ dimensions, as well as the reconstruction map, as shown in figure 14.

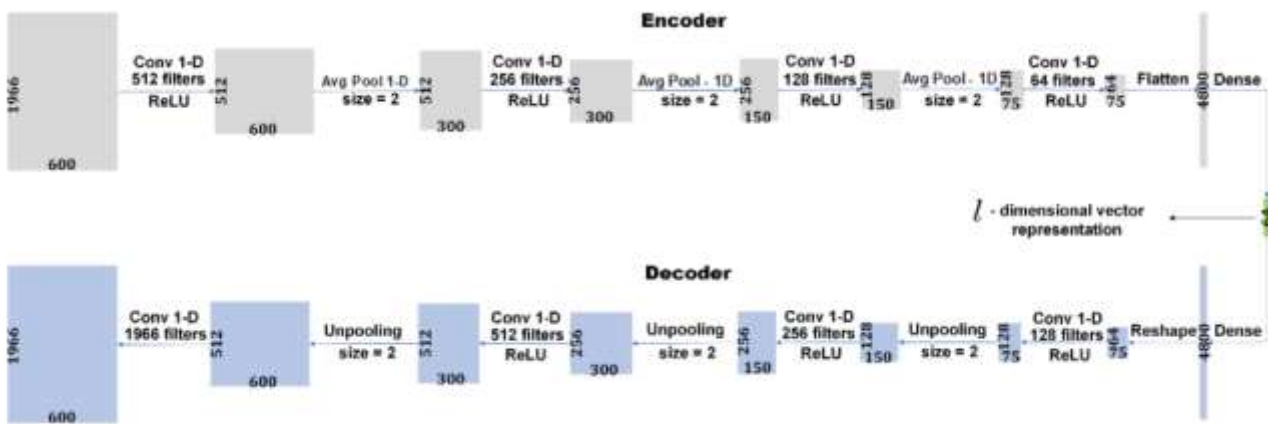


Fig. 14: CAE architecture for the solution of the structural dynamics problem

Then, a feedforward neural network is trained to learn the mapping from the 3-dimensional parametric space to the 64-dimensional encoded space, according to the methodology described in section 3.3. This is the last ingredient needed for the development of the surrogate model, which is then deployed in order to perform a Monte Carlo simulation using an additional triplet of 3000 parameter realizations, i.e. $\{[E_1^i, E_2^i, E_3^i]\}_{i=1}^{3000}$, with minimal computational requirements. The accuracy of the proposed surrogate is validated against the 'exact' solutions of the problem and it is depicted in the following figures 15-19.




Fig. 15: u_x at $t = 4.00\text{sec}$ predicted by (a) the exact model and (b) the surrogate model




Fig. 16: u_y at $t = 4.00\text{sec}$ predicted by (a) the exact model and (b) the surrogate model

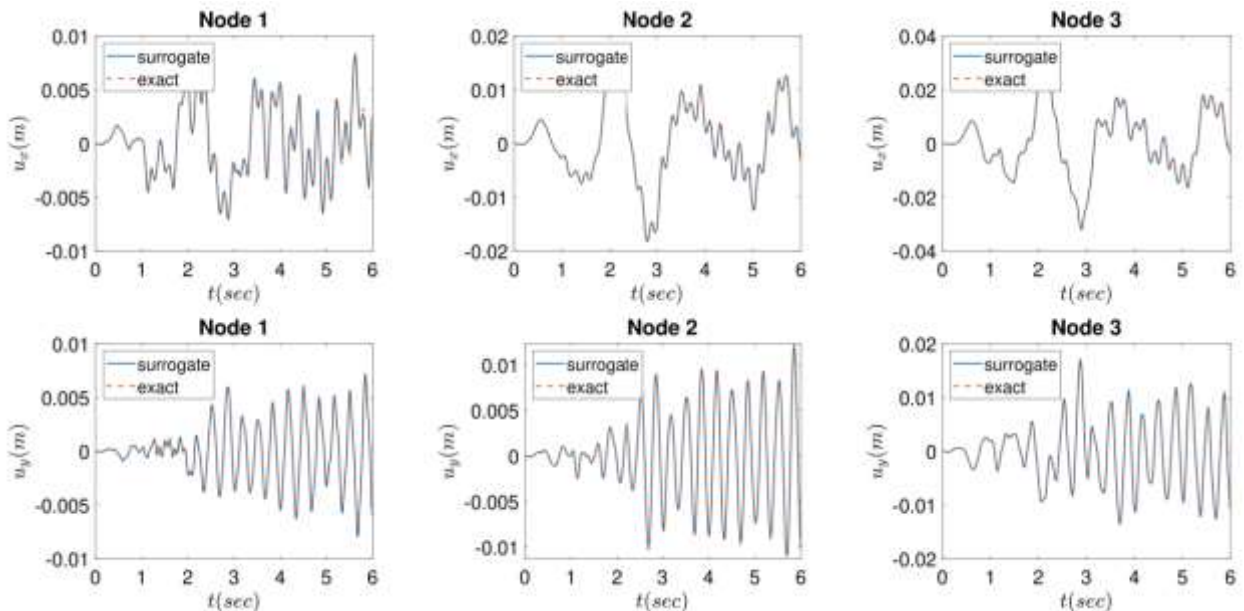


Fig. 17: Displacements u_x and u_y of monitored nodes predicted by the exact and the surrogate model for a specific 'unseen' triplet of parameters $E_1 = 18.66, E_2 = 27.02, E_3 = 21.65 \text{ GPa}$




Fig. 18: Mean value of u_x at $t = 4.00 \text{ sec}$ predicted by (a) the exact and (b) the surrogate model




Fig. 19: Variance of u_x at $t = 4.00 \text{ sec}$ predicted by (a) the exact and (b) the surrogate model

The error in the predictions between the surrogate and the exact (high-fidelity) model was measured to be less than 5% in all cases, while a remarkable decrease in the computational cost for performing the Monte Carlo simulation was achieved, as shown in figure 20(a). The computational gain would have been significantly greater if instead of 3000 simulations we were to perform 50000 simulations (see figure 20(b)), which is a reasonable number in the context of Monte Carlo methods.

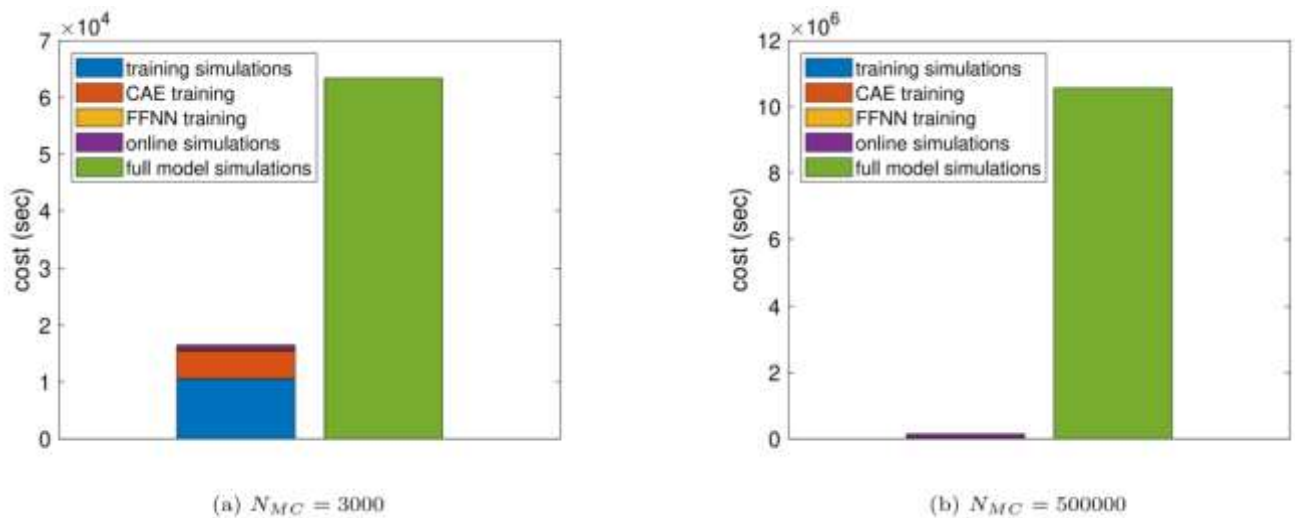


Fig. 20: Comparison of computational cost between the surrogate and the exact model