

POD-DEIM model order reduction for strain softening viscoplasticity

Ghavamian, Fariborz; Tiso, Paolo; Simone, Angelo

Publication date

2017

Document Version

Accepted author manuscript

Published in

Computer Methods in Applied Mechanics and Engineering

Citation (APA)

Ghavamian, F., Tiso, P., & Simone, A. (2017). POD-DEIM model order reduction for strain softening viscoplasticity. *Computer Methods in Applied Mechanics and Engineering*, 317, 458-479.

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

POD-DEIM model order reduction for strain softening viscoplasticity

F. Ghavamian ^a, P. Tiso ^b & A. Simone ^a

^a Faculty of Civil Engineering and Geosciences, Delft University of Technology
P.O. Box 5048, 2600 GA Delft, The Netherlands

^b Institute for Mechanical Systems, ETH Zürich
Leonhardstrasse 21, 8092, Zürich, Switzerland

DRAFT December 5, 2016

Abstract

We demonstrate a Model Order Reduction technique for a system of nonlinear equations arising from the Finite Element Method (FEM) discretization of the three-dimensional quasistatic equilibrium equation equipped with a Perzyna viscoplasticity constitutive model. The procedure employs the Proper Orthogonal Decomposition-Galerkin (POD-G) in conjunction with the Discrete Empirical Interpolation Method (DEIM). For this purpose, we collect samples from a standard full order FEM analysis in the offline phase and cluster them using a novel k -means clustering algorithm. The POD and the DEIM algorithms are then employed to construct a corresponding reduced order model. In the online phase, a sample from the current state of the system is passed, at each time step, to a nearest neighbor classifier in which the cluster that best describes it is identified. The force vector and its derivative with respect to the displacement vector are approximated using DEIM, and the system of nonlinear equations is projected onto a lower dimensional subspace using the POD-G. The constructed reduced order model is applied to two typical solid mechanics problems showing strain localization (a tensile bar and a wall under compression) and a three-dimensional square-footing problem.

KEY WORDS: Finite Element Method, Model Order Reduction, Proper Orthogonal Decomposition, Discrete Empirical Interpolation Method, Perzyna viscoplasticity, strain softening, machine learning, k -means clustering algorithm

1 Introduction

In spite of constant advances in theoretical development and computing power, there are regularly applied many-query procedures, such as parameter sensitivity analysis or parameter estimation for inverse problems, that require many computationally expensive simulations. Model Order Reduction (MOR) techniques aim to reduce the computational cost of these procedures by breaking them into a two-phase procedure commonly known as the offline-online decomposition. A reduced order model is first trained in the computationally expensive offline phase and then used in the online phase to efficiently carry out a many-query procedure. Successful examples can be found in structural dynamics [5], computational fluid dynamics [26, 32] and control systems [20]. Here we employ a set of machine learning techniques that can significantly improve the computational efficiency of a reduced order model.

The main hypothesis of any MOR technique can be put forth as follows [10]:

“For a particular system, the solution space is often attracted to a low-dimensional manifold.”

This hypothesis being true, it is in principle possible to build a set of basis vectors for a lower-dimensional subspace of the solution space, in which solution vectors can be approximated with acceptable accuracy, and therefore reduce the

¹Corresponding author: F. Ghavamian
E-mail addresses: f.ghavamian@tudelft.nl / fari.ghavamian@gmail.com, ptiso@ethz.ch, a.simone@tudelft.nl

computational cost. The lower-dimensional subspace is built by means of samples collected from the solution space and the application of a Principal Component Analysis [17].

Many MOR techniques (Proper Orthogonal Decomposition–Galerkin (POD–G) [22], Balanced Truncation [14] and Moment Matching [3] to cite a few examples) have proven effective in reducing the computational cost of a linear model. However, they all fail to effectively reduce the computational cost of a nonlinear model [10], as the cost of the evaluation of the nonlinear contribution scales to the size of the full order system. To overcome this inefficiency, hyper-reduction methods, such as Missing Point Estimation [4], Gappy Proper Orthogonal Decomposition [31] and Discrete Empirical Interpolation Method (DEIM) [10], have been proposed to approximate the nonlinear contributions.

In the field of solid mechanics, MOR techniques such as POD–G [16, 18], Karhunen Loeve Expansion [12, 19], and Proper Generalized Decomposition (PGD) [1, 11] are employed to reduce the dimension of the system of equations. Nevertheless, the complexity of the evaluation of the constitutive equation remains a computational bottleneck. At variance with these models, in the hyper-reduction technique proposed in Reference [27] the integration of the constitutive equation is performed in a reduced integration domain. In Reference [28], an FE efficient unassembled variant of DEIM (UDEIM) is proposed and applied to geometrically nonlinear structural dynamics problems. In [29], UDEIM is further enhanced to reduce the associated offline cost. Apart from PGD, which is an a-priori method, these contributions focus only on the MOR technique itself, presuming that the underlying samples are fully capable of constructing an efficient reduced order model. We show that this is not necessarily always a valid assumption for the class of problems we are tackling here.

In this contribution we demonstrate an application of a MOR technique –POD–G, Section 3.1– in conjunction with a hyper-reduction method –DEIM, Section 3.2– to build a reduced order model for a boundary value problem –quasistatic equilibrium equation– with a nonlinear and path-dependent constitutive law –Perzyna viscoplasticity [25], Section 2. We observe that the singular values of the collected samples, which we wish had a fast decay as a proof of the central hypothesis of the MOR techniques, decrease at a slow rate. As a result, the efficiency of the reduced order model is diminished. To overcome this issue, Peherstorfer et al. [24] and Haasdonk et al. [15] divide the time domain and the parameter space into several local regions and construct a reduced order model for each one of them individually. Following a similar logic, we employ a set of Machine Learning techniques, such as the k -means clustering algorithm, Section 3.3.1, and the nearest neighbor classifier, Section 3.3.2, as suggested in [24]. As novel contributions, (i) we modify the distance definition in the k -means clustering algorithm to increase the consistency of the clustering procedure and, consequently, the effectiveness of DEIM; (ii) we modify the distance metric in the k -means clustering algorithm and the nearest neighbor classifier, which we heuristically show to be vital for a consistent classification; and (iii) we propose a trivial initial clustering in k -means clustering algorithm that exploits the path-dependent characteristic of the problem. The merits of the proposed MOR technique are demonstrated by three typical solid mechanics problems in Section 4.

2 Problem statement

A strain-softening Perzyna viscoplasticity constitutive model is employed as an archetype of a class of constitutive models typically encountered in nonlinear solid mechanics. The three-dimensional quasistatic equilibrium equation

$$\mathbf{L}^T \boldsymbol{\sigma} = \mathbf{q} \quad \text{in } \Omega, \quad (1)$$

expressed using matrix notation, is defined in the body Ω that is bounded by the surface $\Gamma = \Gamma_u \cup \Gamma_t$. A set of appropriate Dirichlet and Neumann boundary conditions is applied on Γ_u and Γ_t , respectively. In (1), $\boldsymbol{\sigma}$ is an array containing the components of the stress tensor in engineering notation, \mathbf{q} the body force vector, \mathbf{u} the displacement vector, and the operator matrix \mathbf{L} is defined as

$$\mathbf{L}^T = \begin{bmatrix} \partial_x & 0 & 0 & \partial_y & 0 & \partial_z \\ 0 & \partial_y & 0 & \partial_x & \partial_z & 0 \\ 0 & 0 & \partial_z & 0 & \partial_y & \partial_x \end{bmatrix}, \quad (2)$$

in which $\partial_{\#} = \frac{\partial}{\partial \#}$ is the partial derivative operator with respect to $\#$. The equilibrium equation is equipped with the stress-strain relationship

$$\dot{\boldsymbol{\sigma}} = \mathbf{D}^e (\dot{\boldsymbol{\epsilon}} - \dot{\boldsymbol{\epsilon}}^{vp}) \quad (3)$$

in rate form, where $\dot{\#} = \frac{\partial \#}{\partial t}$ is the time derivative of $\#$, \mathbf{D}^e represents the elastic modulus tensor, $\boldsymbol{\varepsilon}$ represents the strain tensor, and $\boldsymbol{\varepsilon}^{\text{vp}}$ the viscoplastic strain tensor. In a Perzyna viscoplasticity constitutive model [25], the viscoplastic strain evolves following the flow rule

$$\dot{\boldsymbol{\varepsilon}}^{\text{vp}} = \dot{\lambda} \frac{\partial \boldsymbol{\theta}}{\partial \boldsymbol{\sigma}} \quad (4)$$

with the rate of the plastic multiplier

$$\dot{\lambda} = \eta \langle \phi_{(\boldsymbol{\theta})} \rangle^{\beta}, \quad (5)$$

the yield function

$$\boldsymbol{\theta} = \boldsymbol{\sigma}_{\text{vm}} - \boldsymbol{\sigma}_{\text{Y}}, \quad (6)$$

and the overstress function

$$\phi_{(\boldsymbol{\theta})} = \frac{\boldsymbol{\theta}}{\boldsymbol{\sigma}_{\text{Y}}}. \quad (7)$$

In the above equations, $\langle \# \rangle$ is the Macaulay bracket, η is the viscosity parameter, β is a model parameter, $\boldsymbol{\sigma}_{\text{Y}}$ is the current yield stress, and $\boldsymbol{\sigma}_{\text{vm}}$ is the von Mises stress. To introduce a strain-softening response, we employ the relation

$$\boldsymbol{\sigma}_{\text{Y}} = \boldsymbol{\sigma}_{\text{Y}_0} ((1+a) e^{-b \kappa} - a e^{-2b \kappa}), \quad (8)$$

where a and b are model parameters, $\boldsymbol{\sigma}_{\text{Y}_0}$ is the initial yield stress, and the plastic strain

$$\kappa = \lambda. \quad (9)$$

The strong form of the quasistatic equilibrium equation (1) is cast into its corresponding weak form and the displacement field is discretized in space using the Finite Element Method (FEM). Standard procedures yield the nonlinear discrete set of equations

$$\mathbf{r}(\mathbf{a}) = \mathbf{f}_{\text{int}}(\mathbf{a}) - \mathbf{f}_{\text{ext}} = \mathbf{0}, \quad (10)$$

where

$$\begin{aligned} \mathbf{f}_{\text{int}}(\mathbf{a}) &= \int_{\Omega} \mathbf{B}^T \boldsymbol{\sigma}(\mathbf{a}) \, d\Omega, \quad \text{and} \\ \mathbf{f}_{\text{ext}} &= \int_{\Gamma_t} \mathbf{N}^T \bar{\mathbf{t}} \, d\Gamma + \int_{\Omega} \mathbf{N}^T \mathbf{q} \, d\Omega. \end{aligned} \quad (11)$$

In (11), $\mathbf{N} \in \mathbb{R}^{3 \times N}$ is a matrix that contains shape functions, $\mathbf{B} = \mathbf{L} \mathbf{N} \in \mathbb{R}^{6 \times N}$ contains derivatives of the shape functions, $\bar{\mathbf{t}}$ are the prescribed tractions on Γ_t , $\mathbf{a} \in \mathbb{R}^N$ is a vector of the nodal values of the displacement vector \mathbf{u} , and N is the total number of degrees of freedom (DOFs).

The system of equations (10) is then solved using a Newton-Raphson scheme according to which

$$\begin{aligned} \mathbf{K}^j \, d\mathbf{a}^{j+1} &= -\mathbf{r}^j, \quad \text{and} \\ \mathbf{a}^{j+1} &= \mathbf{a}^j + d\mathbf{a}^{j+1}, \end{aligned} \quad (12)$$

where $\mathbf{r}^j \in \mathbb{R}^N$ is the residual vector at iteration j , $\mathbf{K}^j \in \mathbb{R}^{N \times N}$ is the stiffness matrix at iteration j , and $d\mathbf{a}^{j+1} \in \mathbb{R}^N$ is the increment of the displacement vector at iteration $j+1$. From here onward, the iteration number j is dropped for the sake of readability. The residual vector

$$\mathbf{r} = \mathbf{K}^L \mathbf{a} + \mathbf{f} - \mathbf{f}_{\text{ext}} \quad (13)$$

consists of a nonlinear part, \mathbf{f} , and a linear part, $\mathbf{K}^L \mathbf{a} - \mathbf{f}_{\text{ext}}$. In the linear part, the matrix

$$\mathbf{K}^L = \int_{\Omega} \mathbf{B}^T \mathbf{D}^e \mathbf{B} \, d\Omega \in \mathbb{R}^{N \times N} \quad (14)$$

is the linear elastic stiffness matrix, equal to the stiffness matrix \mathbf{K} at the first iteration of the first time step. In the nonlinear part, the vector $\mathbf{f} = \mathbf{f}(\mathbf{a}) \in \mathbb{R}^N$ is referred to as the force vector. The derivative of the force vector, $\frac{\partial \mathbf{f}}{\partial \mathbf{a}} \in \mathbb{R}^{N \times N}$, generates the nonlinear part of the stiffness matrix.

Strain softening viscoplasticity typically induces strongly localized strain fields. As a result, a large number of degrees of freedom is required to resolve these gradients, thus causing two computational bottlenecks:

1. the solution of an $N \times N$ system (12.1), and
2. the evaluation of the force vector \mathbf{f} and its derivative with respect to the displacement vector $\frac{\partial \mathbf{f}}{\partial \mathbf{a}}$.

3 Model Order Reduction

Essentially, a Model Order Reduction technique decomposes a many-query procedure into two parts: a computationally expensive offline phase, and a computationally efficient online phase. In the offline phase, a set of samples is collected from a standard analysis (in this context, this is done using FEM). This information is employed to construct a reduced order model. In the online phase, the reduced order model is used to carry out the many-query procedure. The MOR technique employed in this study is summarized in Section 3.5 while the corresponding online/offline procedures are detailed next.

The POD-G [10] is employed to project the nonlinear system of equations (10) onto a lower-dimensional subspace. DEIM [10] is then used to approximate the force vector and its derivative with respect to the displacement vector. In practice, the Proper Orthogonal Decomposition [8] is applied to a set of samples collected from a full order FEM analysis to compute a set of basis vectors. These basis vectors are later used in the POD-G and the DEIM procedures. Finally, the Localized Discrete Empirical Interpolation Method (LDEIM) [24] is employed to overcome a shortcoming of DEIM that occurs when the collected samples of the force vector are not attracted to a sufficiently lower-dimensional subspace. At variance with [24], we propose (i) a modified distance definition in the k -means clustering algorithm that is tailored for DEIM, (ii) a modified distance metric that is used in the k -means clustering algorithm and the nearest neighbor classifier, and (iii) a trivial initial clustering in the k -means clustering algorithm that exploits the path-dependent characteristic of the problem.

3.1 Proper Orthogonal Decomposition-Galerkin

3.1.1 Galerkin projection

To begin with, we attempt to project the nonlinear system of equations (10) onto a lower-dimensional subspace. Suppose that the solution of (10) is attracted to a lower-dimensional subspace $\mathcal{Y} \subset \mathbb{R}^N$. This solution can then be approximated by a weighted linear combination of the set of orthonormal basis vectors $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ ($k \ll N$) spanning \mathcal{Y} as in

$$\mathbf{a} \approx \mathbf{V} \tilde{\mathbf{a}}, \quad (15)$$

where $\tilde{\mathbf{a}} \in \mathbb{R}^k$ is a set of scalar weights associated to the set of basis vectors $\mathbf{V} \in \mathbb{R}^{N \times k}$. Inserting (15) into (10) yields an overdetermined nonlinear system of N equations in k unknowns. To compute an approximate value for $\tilde{\mathbf{a}}$, we construct an equivalent system of equations $\mathbf{g} = \mathbf{0}$ that is determined. To find such a system, we search for a set of k functions \mathbf{g} in subspace \mathcal{Y} such that, for any given value of $\tilde{\mathbf{a}}$, the L_2 norm of the difference between the residual $\mathbf{r}(\mathbf{V}\tilde{\mathbf{a}})$ and its counterpart in the lower-dimensional subspace $\mathbf{V}\mathbf{g}$ is minimized according to [7, Theorem 3.2]

$$\arg \min_{\mathbf{g}} \|\mathbf{r}(\mathbf{V}\tilde{\mathbf{a}}) - \mathbf{V}\mathbf{g}\|_2. \quad (16)$$

Owing to the choice of the L_2 norm, it is straightforward to solve (16) for \mathbf{g} . Consequently, the equivalent system of equations $\mathbf{g} = \mathbf{0}$ becomes

$$\mathbf{V}^T \mathbf{r}(\mathbf{V}\tilde{\mathbf{a}}) = \mathbf{0}. \quad (17)$$

The system of equations (17) is then solved using a Newton-Raphson scheme according to which

$$\begin{aligned} \tilde{\mathbf{K}}^j d\tilde{\mathbf{a}}^{j+1} &= -\tilde{\mathbf{r}}^j, & \text{and} \\ \tilde{\mathbf{a}}^{j+1} &= \tilde{\mathbf{a}}^j + d\tilde{\mathbf{a}}^{j+1}, \end{aligned} \quad (18)$$

where $\tilde{\mathbf{K}}^j = \mathbf{V}^T \mathbf{K}^j \mathbf{V} \in \mathbb{R}^{k \times k}$ is the reduced stiffness matrix and $\tilde{\mathbf{r}}^j = \mathbf{V}^T \mathbf{r}^j \in \mathbb{R}^k$ is the reduced residual vector. The set of basis vectors \mathbf{V} is computed using the Proper Orthogonal Decomposition as explained in Section 3.1.2.

In spite of the system of equations (18) being of dimension $k \ll N$, the complexity of problem still depends on N when computing the nonlinear contributions. Efficient reduction is achieved later in Sections 3.2 and 3.3 by approximating the force vector and its derivative with respect to the displacement vector using DEIM.

3.1.2 Proper Orthogonal Decomposition

A set of orthonormal basis vector $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ is used to project the system of equations (10) onto a lower-dimensional subspace as in (17). This set is computed by means of the Proper Orthogonal Decomposition (POD) method [8, 30].

For this purpose, samples of the displacement and the force vectors are collected from a standard full order FEM analysis and gathered in the sample matrices $\mathbf{A}^{\text{samp}} = [\mathbf{a}_1, \dots, \mathbf{a}_{n_s}]$ and $\mathbf{F}^{\text{samp}} = [\mathbf{f}_1, \dots, \mathbf{f}_{n_s}]$, respectively. For the displacement samples (and similarly for the force samples), we search for a subspace $\mathcal{Y} = \text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_k\} \subset \mathbb{R}^N$ in which the samples can be optimally described. In other words, it is preferable that the error between samples and their projection on \mathcal{Y} is minimized in a certain sense. Formally, the set of orthonormal basis vectors spanning \mathcal{Y} is computed by solving the optimization problem [10]

$$\begin{aligned} & \arg \min_{\{\mathbf{v}_1, \dots, \mathbf{v}_k\}} \sum_{i=1}^{n_s} \|\mathbf{a}_i - \sum_{j=1}^k (\mathbf{v}_j^T \mathbf{a}_i) \mathbf{v}_j\|_2^2, \\ & \text{subjected to } \mathbf{v}_m^T \mathbf{v}_n = \begin{cases} 1 & m = n \\ 0 & m \neq n \end{cases}, \quad m, n = 1, \dots, k. \end{aligned} \quad (19)$$

In the equation above, the error between a sample \mathbf{a}_i and its projection $\sum_{j=1}^k (\mathbf{v}_j^T \mathbf{a}_i) \mathbf{v}_j$ onto the subspace \mathcal{Y} is minimized in an L_2 sense. Note that the norms in (19) are squared to facilitate the derivation of (21). It can be shown [30] that the solution to (19) is a set of left singular vectors \mathbf{V} obtained from the application of the Singular Value Decomposition (SVD) to the sample matrix \mathbf{A}^{samp} such that

$$\mathbf{V} \mathbf{\Sigma} \mathbf{W}^T = \mathbf{A}^{\text{samp}}, \quad (20)$$

where $\mathbf{V} \in \mathbb{R}^{N \times N}$ is a set of left orthonormal singular basis vectors, $\mathbf{W}^{n_s \times n_s}$ is a set of right orthonormal singular basis vectors, and $\mathbf{\Sigma}$ is a diagonal matrix that contains the singular values in a descending manner $\sigma_1 > \dots > \sigma_r$, $r = \text{rank}(\mathbf{A}^{\text{samp}})$. Therefore, by using an SVD [30, Theorem 1.1.1], the residual of (19) becomes [10]

$$\sum_{i=1}^{n_s} \|\mathbf{a}_i - \sum_{j=1}^k (\mathbf{v}_j^T \mathbf{a}_i) \mathbf{v}_j\|_2^2 = \sum_{i=k+1}^r \sigma_i^2. \quad (21)$$

From the equation above, it is possible to infer that the projection error scales to the sum of the squares of the neglected ($i = k + 1$ to r) singular values (RHS of (21)). As a result of this observation, a heuristic criterion to select the number of basis vectors k would be to consider the ratio of the sum of the squares of the neglected singular values to the sum of all the singular values as in

$$\frac{\sum_{i=k+1}^r \sigma_i^2}{\sum_{i=1}^r \sigma_i^2} \leq \varepsilon \quad (22)$$

where $\varepsilon \in [0, 1]$ is a user-defined tolerance. As the value of ε deviates from one and tends to zero, the number of neglected singular values decreases. Simultaneously, the number of basis vectors which are taken into account increases, yielding a more accurate reduced order model. The POD procedure is summarized in Algorithm 1.

Algorithm 1: POD procedure

input : A set of sample \mathbf{A}^{samp} , a tolerance ε

output: A set of truncated basis vectors $\mathbf{V} \in \mathbb{R}^{N \times k}$

Apply Singular Value Decomposition: $\mathbf{A}^{\text{samp}} = \mathbf{V} \mathbf{\Sigma} \mathbf{W}^T$

Truncate \mathbf{V} based on $\mathbf{\Sigma}$ and criterion (22)

3.2 Discrete Empirical Interpolation Method

We now tackle the computational bottleneck related to the complexity of the evaluation of the force vector and its derivative with respect to the displacement vector. This is done by employing DEIM [10]. Suppose that the force vector \mathbf{f} in (13) is attracted to a lower-dimensional subspace $\mathcal{Z} = \text{span}\{\mathbf{u}_1, \dots, \mathbf{u}_m\} \subset \mathbb{R}^N$, where $m \ll N$. The force vector can then be approximated using a weighted linear combination of the basis vectors spanning \mathcal{Z} according to

$$\mathbf{f} \approx \hat{\mathbf{f}} = \mathbf{U} \mathbf{c}, \quad (23)$$

where $\mathbf{c} \in \mathbb{R}^m$ is a set of scalar weights associated to the set of basis vectors \mathbf{U} . To find these scalar weights, m rows of both sides of (23) are selected by means of a matrix $\mathbf{P} \in \mathbb{R}^{N \times m}$:

$$\mathbf{P}^T \mathbf{f} \approx \mathbf{P}^T \mathbf{U} \mathbf{c}. \quad (24)$$

Solving (24) for \mathbf{c} and inserting it in (23) yield the approximated force vector

$$\hat{\mathbf{f}} = \mathbf{U} (\mathbf{P}^T \mathbf{U})^{-1} \mathbf{P}^T \mathbf{f} = \mathbf{U} (\mathbf{P}^T \mathbf{U})^{-1} \mathbf{f}_{\text{DEIM}}. \quad (25)$$

Here, $\mathbf{f}_{\text{DEIM}} = \mathbf{P}^T \mathbf{f} \in \mathbb{R}^m$ contains m entries of the force vector, and the matrix \mathbf{P} is defined with the aid of the DEIM algorithm listed in the next subsection. The complexity of the evaluation of $\hat{\mathbf{f}}$ is now reduced to that of $\hat{\mathbf{f}}$; thus, this consists of the evaluation of \mathbf{f}_{DEIM} and a matrix multiplication (note that the computation of $\mathbf{U} (\mathbf{P}^T \mathbf{U})^{-1}$ is carried out in the offline phase).

To compute the derivative of the approximated force vector (25) with respect to the displacement vector, several methods are suggested in [6]. Here we choose to utilize the same set of basis vectors previously computed for the force vector, as originally suggested in [10]:

$$\frac{\partial \hat{\mathbf{f}}}{\partial \mathbf{v} \bar{\mathbf{a}}} = \mathbf{U} (\mathbf{P}^T \mathbf{U})^{-1} \mathbf{P}^T \frac{\partial \mathbf{f}}{\partial \mathbf{v} \bar{\mathbf{a}}} = \mathbf{U} (\mathbf{P}^T \mathbf{U})^{-1} \mathbf{K}_{\text{DEIM}}, \quad (26)$$

where $\mathbf{K}_{\text{DEIM}} = \mathbf{P}^T \frac{\partial \mathbf{f}}{\partial \mathbf{v} \bar{\mathbf{a}}} \in \mathbb{R}^{m \times N}$ contains m rows of the nonlinear part of the stiffness matrix. The procedure to assemble \mathbf{f}_{DEIM} and \mathbf{K}_{DEIM} is detailed in Algorithm 2.

Algorithm 2: Assembly of \mathbf{K}_{DEIM} and \mathbf{f}_{DEIM}

input : A set $\mathbf{IND} = [\mathbf{IND}_1, \dots, \mathbf{IND}_m]$ of entries selected by the DEIM Algorithm 3

output: \mathbf{K}_{DEIM} and \mathbf{f}_{DEIM}

Set \mathbf{K}_{DEIM} and \mathbf{f}_{DEIM} to zero

for $i = 1, \dots, m$ **do**

Find elements $\mathbf{E} = \{E_1, \dots, E_{n_e}\}$ sharing the \mathbf{IND}_i entry

for $e = 1 \dots n_e$ **do**

Compute the nonlinear part of element stiffness matrix and the force vector: $\mathbf{K}_e^{\mathbf{E}}$ and $\mathbf{f}_e^{\mathbf{E}}$

Find the local DOF j associated to the entry \mathbf{IND}_i

$\mathbf{f}_{\text{DEIM}}(i) = \mathbf{f}_{\text{DEIM}}(i) + \mathbf{f}_e^{\mathbf{E}}(j)$

Define \mathbf{gIndex} as the vector of global DOFs associated with element e

$\mathbf{K}_{\text{DEIM}}(i, \mathbf{gIndex}) = \mathbf{K}_{\text{DEIM}}(i, \mathbf{gIndex}) + \mathbf{K}_e^{\mathbf{E}}(j, :)$

end

end

3.2.1 Discrete Empirical Interpolation Method algorithm

The error bound for the DEIM approximation $\hat{\mathbf{f}}$ of the force vector \mathbf{f} is given by [10]

$$\|\mathbf{f} - \hat{\mathbf{f}}\|_2 \leq \|(\mathbf{P}^T \mathbf{U})^{-1}\|_2 \|(\mathbf{I} - \mathbf{U} \mathbf{U}^T) \mathbf{f}\|_2. \quad (27)$$

To increase the accuracy of the DEIM approximation, the value of $\|(\mathbf{I} - \mathbf{U} \mathbf{U}^T) \mathbf{f}\|_2$ should be decreased while the value of $\|(\mathbf{P}^T \mathbf{U})^{-1}\|_2$ is kept bounded [10, Lemma 3.2].

The quantity $\|(\mathbf{I} - \mathbf{U} \mathbf{U}^T) \mathbf{f}\|_2$ is the projection error induced by POD (21). Therefore, by increasing the number of basis vectors in \mathbf{U} , according to (22), its value becomes smaller. However, by increasing the number of basis vectors, the value of $\|(\mathbf{P}^T \mathbf{U})^{-1}\|_2$ becomes larger, jeopardizing the accuracy of the DEIM approximation. According to Chaturantabut and Sorensen [10, Lemma 3.2], the DEIM Algorithm 3 computes the matrix \mathbf{P} in a way that the growth of this term is bounded. Consequently, the accuracy of the approximation can be improved by taking more and more basis vectors into account.

In Algorithm 3, \mathbf{e}_i is the i th column of an identity matrix. The notation $\mathbf{M} \leftarrow [\mathbf{M}, \mathbf{v}_i]$ indicates that the matrix \mathbf{M} is augmented by adding the column vector \mathbf{v}_i .

Algorithm 3: DEIM algorithm

input : A set of basis vectors $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_m]$
output: Entries of the force vector selected by DEIM $\mathbf{IND} = [IND_1, \dots, IND_m]$, and the matrix $\mathbf{P} = [\mathbf{e}_{IND_1}, \dots, \mathbf{e}_{IND_m}]$

Set $IND_1 =$ index of the maximum entry of $|\mathbf{u}_1|$
Initiate matrices: $\mathbf{W} = [\mathbf{u}_1]$, $\mathbf{P} = [\mathbf{e}_{IND_1}]$, $\mathbf{IND} = [IND_1]$
for $i = 2, \dots, m$ **do**
 Solve $\mathbf{P}^T \mathbf{W} \mathbf{c} = \mathbf{P}^T \mathbf{u}_i$ for \mathbf{c}
 Compute residual $\mathbf{r} = \mathbf{u}_i - \mathbf{W} \mathbf{c}$
 Set $IND_i =$ index of the maximum component of $|\mathbf{r}|$
 Augment $\mathbf{W} \leftarrow [\mathbf{W}, \mathbf{u}_i]$, $\mathbf{P} \leftarrow [\mathbf{P}, \mathbf{e}_{IND_i}]$ and $\mathbf{IND} \leftarrow [\mathbf{IND}, IND_i]$
end

3.3 Localized Discrete Empirical Interpolation Method

DEIM succeeds in defining an efficient reduced order model, as in [10], only if the singular values of the force vector decay sufficiently fast –this effectively means that, when a relatively small number of basis vectors are considered, the sum of the squares of the singular values corresponding to the neglected basis vectors becomes negligibly small. However, in the context of a viscoplasticity constitutive model this assumption is, apparently, not valid. As it will be shown in Section 4.2, a considerably large number of basis vectors are required to sufficiently reduce the projection error. Hence, an efficient reduced order model cannot be constructed merely by applying the standard version of DEIM.

We employ a Localized DEIM [24] to remedy the above shortcoming. The general idea is to group samples, which are close to each other in a certain sense, into clusters. We show that the decay of the singular values becomes significantly faster for each cluster, implying that the samples are attracted to a considerably lower-dimensional subspace. In this work we only investigate the performance of DEIM in conjunction with the clustered force samples. Nonetheless, the displacement samples could also be clustered as suggested in [6]. For the purpose of clustering the samples, we make use of machine learning techniques such as the k -means clustering algorithm (for clustering samples in the offline phase), and the nearest neighbor classifier (for finding the correct cluster in the online phase).

3.3.1 Modified k -means clustering algorithm

We now describe a procedure to divide a set of samples \mathbf{F}^{samp} , schematically shown in Figure 1a, into n_c clusters with the intention that samples in each cluster are close to each other according to a certain distance definition. For this purpose, a k -means clustering algorithm, informally explained in [23], is used. In this contribution the k -means clustering algorithm is tailored for DEIM and the path-dependent nature of the problem. More specifically, we will improve on the cluster initiation, the distance definition, and the distance metric.

The strong path-dependency of the problem under consideration leads to a rather straightforward choice of the initial guess for the clustering. In other words, since force vectors from a certain part of the loading history are already “close” to each other, the samples are simply divided into n_c clusters without any permutation as an initial guess for the k -means clustering algorithm.

In contrast to the definition of the distance in the standard k -means clustering algorithm, which is the distance between a sample and the center of clusters, the distance is here defined between a sample and its DEIM approximation. For instance, consider a sample \mathbf{f}_i and, for each cluster $c = 1 \dots n_c$, a set of basis vectors ${}_c \mathbf{U}$ and a matrix ${}_c \mathbf{P}$ (from here onward, the number of clusters is indicated by a subscript on the left-hand side of the corresponding quantity). The distance measure used in the k -means clustering algorithm is defined by the difference between \mathbf{f}_i and its DEIM approximation ${}_c \hat{\mathbf{f}}_i = {}_c \mathbf{A} {}_c \mathbf{P}^T \mathbf{f}_i$ for each cluster $c = 1 \dots n_c$, where ${}_c \mathbf{A} = {}_c \mathbf{U} ({}_c \mathbf{P}^T {}_c \mathbf{U})^{-1}$. In the standard k -means clustering algorithm, it is typical to consider a Euclidean distance metric. Shortcomings of the Euclidean distance are further discussed in [2]. Here, a stricter metric, based on the Euclidean distance and the cosine similarity, is defined:

$$d = \sqrt{d_{\text{Euclidean}}^2 + d_{\text{cosine}}^2}, \quad (28)$$

where

$$d_{\text{Euclidean}} = \|\mathbf{f}_i - {}_c\hat{\mathbf{f}}_i\|_2 / \|\mathbf{f}_i\|_2 \quad \text{and} \quad d_{\text{cosine}} = 1 - |\mathbf{f}_i^T {}_c\hat{\mathbf{f}}_i| / (\|\mathbf{f}_i\|_2 \|\mathbf{f}_i\|_2). \quad (29)$$

Since $\|\mathbf{f}_i - {}_c\hat{\mathbf{f}}_i\|_2 < \|\mathbf{f}_i\|_2$, these two measures are comparable in terms of magnitude; values of $d_{\text{Euclidean}}$ and d_{cosine} fall between zeros and one, and it is therefore safe to take their norm as in (28).

The proposed k -means clustering algorithm is detailed in Algorithm 4. The procedure initiates by dividing samples in \mathbf{F}^{samp} into n_c clusters, as schematically shown in Figure 1b, without permuting them. Additionally, a mapping (**sampleToClus** : sample \rightarrow cluster number) is initiated to keep track of samples and their corresponding cluster number. In the k -means iterative scheme, the POD (Algorithm 1) and the DEIM (Algorithm 3) procedures are applied to each one of the clusters, and a reduced set of basis vectors ${}_c\mathbf{U}$, a matrix ${}_c\mathbf{P}$, and a set of DEIM selected entries ${}_c\mathbf{IND}$ are computed. The distance between each sample \mathbf{f}_i and its DEIM approximation ${}_c\hat{\mathbf{f}}_i = {}_c\mathbf{A} {}_c\mathbf{P}^T \mathbf{f}_i$ (28) is saved in a matrix $\mathbf{G} \in \mathbb{N}^{n_s \times n_c}$ (n_s being the number of samples).

Next, each sample \mathbf{f}_i is checked to make sure that it belongs to its cluster. To do so, ${}_c\hat{\mathbf{f}}_i = {}_c\mathbf{A} {}_c\mathbf{P}^T \mathbf{f}_i$ is computed for each one of the clusters $c = 1 \dots n_c$. Then, the distance between \mathbf{f}_i and each one of the ${}_c\hat{\mathbf{f}}_i$ is measured by means of (28). The cluster that helps generating the most accurate DEIM approximation of \mathbf{f}_i , such that ${}_c\hat{\mathbf{f}}_i$ is closest to \mathbf{f}_i as shown in Figure 1c, is the cluster \mathbf{f}_i actually belongs to. To keep track of samples in the clusters, when a sample is designated to another cluster, the mapping **sampleToClus** is updated. The procedure continues until a maximum number of iterations is reached or all the samples are (sub)optimally clustered.

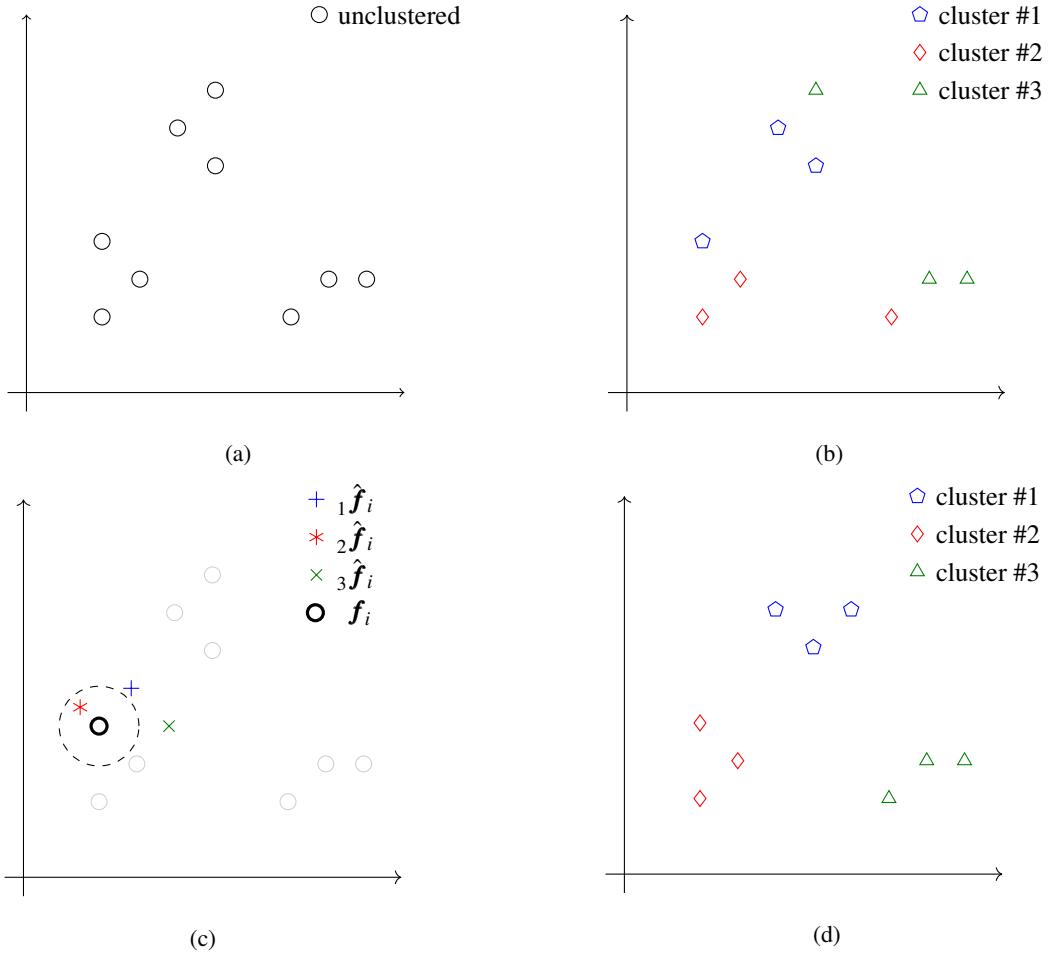


Figure 1: Modified k -means clustering algorithm: (a) unclustered samples; (b) initial clustering of samples; (c) distance between a sample \mathbf{f}_i and its DEIM approximation using ${}_c\mathbf{A}$ and ${}_c\mathbf{P}$ of all clusters, $c = 1 \dots n_c$. The sample \mathbf{f}_i belongs to the cluster that generates the closest ${}_c\hat{\mathbf{f}}_i = {}_c\mathbf{A} {}_c\mathbf{P}^T \mathbf{f}_i$ in the sense of (28); (d) final form of clusters after convergence of the modified k -means clustering algorithm.

When an insufficient number of basis vectors is considered, the rate of convergence of the k -means clustering algorithm becomes very slow and the quality of the clusters diminishes. From the implementation point of view, to prevent division by zero in d_{cosine} and $d_{\text{Euclidean}}$ (29), it is vital to eliminate zero vectors from the matrix of force samples.

In order to check the validity of the proposed approach we can plot the singular values associated with each cluster and observe their rate of decay. If their rate of decay is not substantially faster than that of the singular values associated with all the samples, the number of clusters needs to be increased at the cost of a more expensive k -means clustering algorithm.

Algorithm 4: Modified k -means clustering algorithm	
<p>input : A set of force vector samples $\mathbf{F}^{\text{samp}} = [\mathbf{f}_1, \dots, \mathbf{f}_{n_s}]$, total number of clusters n_c, maximum number of iterations maxIter</p> <p>output: ${}_c\mathbf{A}$, ${}_c\mathbf{P}$, ${}_c\text{IND}$ where $c = 1, \dots, n_c$, the mapping <i>sampleToClus</i></p> <p>Cluster samples without sorting $\mathbf{F}^{\text{samp}} = \{\mathbf{f}_1, \dots, \mathbf{f}_{n_s}\}$ Initiate <i>sampleToClus</i> (mapping from a sample to its corresponding cluster) Set <i>converged</i> = 1 while <i>converged</i> \neq 1 and <i>maxIter</i> not reached do for $c = 1 \dots n_c$ do Compute a set of reduced basis vectors: ${}_c\mathbf{U} = \text{POD}({}_c\mathbf{F})$ Compute the DEIM entries: $[{}_c\text{IND}, {}_c\mathbf{P}] = \text{DEIM}({}_c\mathbf{U})$ Compute the DEIM coefficient matrix: ${}_c\mathbf{A} = {}_c\mathbf{U} ({}_c\mathbf{P}^T {}_c\mathbf{U})^{-1}$ for $i = 1$ to n_s do Compute the DEIM approximation: $\hat{\mathbf{f}}_i = {}_c\mathbf{A} {}_c\mathbf{P}^T \mathbf{f}_i$ Compute the distance d and save it in $\mathbf{G}(i, c)$ end end for $i = 1$ to n_s do Get the old cluster number: $c_{\text{old}} = \text{sampleToClus}(i)$ Find the new cluster number: $c_{\text{new}} = \text{index of the minimum value of } \mathbf{G}(i, :)$ if $c_{\text{old}} \neq c_{\text{new}}$ then <i>converged</i> = 0 Update mapping and, consequently, clusters : $\text{sampleToClus}(i) = c_{\text{new}}$ end end if <i>converged</i> = 1 then Save ${}_c\mathbf{A}$, ${}_c\mathbf{P}$, and ${}_c\text{IND}$ for all $c = 1 \dots n_c$ Save the mapping <i>sampleToClus</i> end end</p>	<p>Algorithm 1</p> <p>Algorithm 3</p> <p>Equation (28)</p>

3.3.2 Nearest neighbor classifier

To find the cluster that best represents the current state of the system during the runtime, we employ a nearest neighbor classifier and modify its distance metric according to (28). At each time step, the nearest neighbor classifier computes distances (28) between the converged force vector $\hat{\mathbf{f}}$ and each one of the clustered samples \mathbf{f}_i , $i = 1 \dots n_s$. The classifier then determines the sample with the least distance to the converged force vector. The cluster that contains this sample best describes the current state of the system.

To reduce the computational cost of the classification, the nearest neighbor classifier uses a low-dimensional representation of the force vector. In this context, we employ the DEIM-based feature extraction [24]

$$\hat{\mathbf{f}}_{\text{FE}} = {}_{c_{\text{old}}}\mathbf{P}^T \hat{\mathbf{f}} \quad \text{and} \quad \mathbf{f}_{\text{FE},i} = {}_{c_{\text{old}}}\mathbf{P}^T \mathbf{f}_i, \quad (30)$$

where $\hat{\mathbf{f}}_{\text{FE}} \in \mathbb{R}^m$ and $\mathbf{f}_{\text{FE},i} \in \mathbb{R}^m$ are a low-dimensional representation of $\hat{\mathbf{f}}$ and \mathbf{f}_i , respectively, and ${}_{c_{\text{old}}}\mathbf{P} \in \mathbb{R}^{N \times m}$ is the DEIM matrix from the previous time step. Distance metric in (28) are computed as

$$d_{\text{Euclidean}} = \|\mathbf{f}_{\text{FE},i} - \hat{\mathbf{f}}_{\text{FE}}\|_2 / \|\mathbf{f}_{\text{FE},i}\|_2 \quad \text{and} \quad d_{\text{cosine}} = 1 - |\mathbf{f}_{\text{FE},i}^T \hat{\mathbf{f}}_{\text{FE}}| / (\|\mathbf{f}_{\text{FE},i}\|_2 \|\hat{\mathbf{f}}_{\text{FE}}\|_2). \quad (31)$$

Algorithm 5: Nearest neighbor classifier

input : A set of force vector samples $\mathbf{F}^{\text{samp}} = [\mathbf{f}_1, \dots, \mathbf{f}_{n_s}]$,
the *sampleToClus* (mapping from a sample to its corresponding cluster),
the force vector $\hat{\mathbf{f}} \in \mathbb{R}^N$ at time step t ,
the current cluster c_{old}

output: New cluster c_{new}

Extract feature of the force vector: $\hat{\mathbf{f}}_{\text{FE}} = {}_{c_{\text{old}}}\mathbf{P}^T \hat{\mathbf{f}}$

for $i = 1 \dots n_s$ **do**

 Extract feature of the sample: $\mathbf{f}_{\text{FE},i} = {}_{c_{\text{old}}}\mathbf{P}^T \mathbf{f}_i$

 Compute the distance $\mathbf{d}(i)$ between $\hat{\mathbf{f}}_{\text{FE}}$ and $\mathbf{f}_{\text{FE},i}$ Equations (28) and (31)

end

Find the index i_{min} of the minimum component of \mathbf{d}

Find the new cluster $c_{\text{new}} = \text{sampleToClus}(i_{\text{min}})$

3.4 Discussion on the computational cost

The computational cost of DEIM and LDEIM is already discussed in [10, 24]. For the sake of completeness, this section is dedicated to a brief review of the key points. Moreover, we also discuss some of our observations.

- In spite of the additional computational cost of the clustering, the cost of the offline phase is still strongly dominated by the sample collection.
- For the online phase, the primary contributor to the computational cost is the nearest neighbor classifier, Algorithm 5. Note that the computational complexity of this procedure is independent of the number of dofs N , thanks to the feature extraction concept (30). However, it strongly depends on the total number of samples n_s . As n_s becomes larger, the nearest neighbor classification loses its efficiency. To tackle this problem, it is possible to either reduce the number of samples through Redundancy Reduction methods such as the QR-SVD technique [9, 21], or perform the classification based on a set of surrogate quantities instead of samples. For instance, the centroid of the clusters computed in classical k-means clustering algorithm can be an adequate substitute. In this manner, the procedure becomes dependent on the number of clusters. However, hopefully, the number of clusters are always significantly smaller than the number of samples.
- If, as a result of the POD and the DEIM approximations, the consistency of the linearization of (17) is jeopardized, the number of iterations in the Newton-Raphson scheme in the reduced order model will increase. Nevertheless, the iteration counts could be kept the same by accepting a solution that is not fully converged. This claim could be conjectured as follows. Consider the error between the full order model and the outcome of the Newton-Raphson scheme in the reduced order model, $\mathbf{e} = \mathbf{a} - \mathbf{a}^{\text{NR}}$. By adding and subtracting the exact solution of (17), \mathbf{a}^{ROM} , to the right hand side of the error expression and taking the L_2 norm of both sides, an upper bound to the error becomes

$$\|\mathbf{e}\|_2 \leq \|\mathbf{a} - \mathbf{a}^{\text{ROM}}\|_2 + \|\mathbf{a}^{\text{ROM}} - \mathbf{a}^{\text{NR}}\|_2. \quad (32)$$

From the equation above, it can be inferred that, regardless of how accurately (17) is solved, even at the limit where $\|\mathbf{a}^{\text{ROM}} - \mathbf{a}^{\text{NR}}\|_2$ becomes effectively zero, the error as a result of the MOR approximation $\|\mathbf{a} - \mathbf{a}^{\text{ROM}}\|_2$ remains untouched. As a result, the bound on the error is, at most, as tight as $\|\mathbf{a} - \mathbf{a}^{\text{ROM}}\|_2$ permits. Hence, a solution that is relatively accurate, considering the error as a result of MOR, but not fully converged in the Newton-Raphson scheme should be considered as acceptable. Nevertheless, in Section 4, we always consider enough basis vectors to ensure the consistency of the linearization and, consequently, quadratic convergence rate of the Newton-Raphson scheme.

- In theory, as the number of clusters becomes larger, the computer memory occupied by the basis vectors in the online phase grows dramatically. However, in practice, we experienced that by increasing the number of clusters, the number of basis vectors in each cluster becomes significantly smaller. As a result, the cumulative number of basis vectors does not increase dramatically, as shown in Section 4.3.

3.5 Summary of the proposed Model Order Reduction technique

Next, a step-by-step guide for the proposed MOR technique is outlined. The offline phase (Algorithm 6) is run first and its output data are saved to a file. In the online phase (Algorithm 7), the data are read and the reduced order model is employed to run the simulation.

Algorithm 6: Offline phase	
1. Collect samples of the displacement and the force vectors in \mathbf{A}^{samp} and \mathbf{F}^{samp} , respectively, using a standard FEM analysis.	
2. Compute a set of POD basis vectors for the displacement samples: $\mathbf{V} = \text{POD}(\mathbf{A}^{\text{samp}})$.	Section 3.1.2
3. Cluster force samples using the k -means clustering algorithm (Algorithm 4), and compute the converged values of the DEIM coefficient matrices ${}_c\mathbf{A}$, the matrix ${}_c\mathbf{P}$, the DEIM selected entries ${}_c\mathbf{IND}$ for each cluster ($c = 1 \dots n_c$), and the mapping <i>sampleToClus</i> . In Algorithm 4 for $c = 1 \dots n_c$:	Section 3.3.1
3.1. use the POD procedure (Algorithm 1): ${}_c\mathbf{U} = \text{POD}({}_c\mathbf{F})$,	Section 3.1.2
3.2. use the DEIM Algorithm 3: $[{}_c\mathbf{IND}, {}_c\mathbf{P}] = \text{DEIM}({}_c\mathbf{U})$,	Section 3.2.1
3.3. compute ${}_c\mathbf{A} = {}_c\mathbf{U} ({}_c\mathbf{P}^T {}_c\mathbf{U})^{-1}$, and	
3.4. update the mapping matrix <i>sampleToClus</i> .	
4. Save $\mathbf{V}, \mathbf{F}^{\text{samp}}, \text{sampleToClus}$, and the converged values of ${}_c\mathbf{A}, {}_c\mathbf{P}, {}_c\mathbf{IND}$.	

4 Applications

We demonstrate the proposed MOR technique by means of three typical solid mechanics problems. A specific set of parameters ($\eta = 10^{-5} \text{ s}^{-1}$, $a = -1$, $b = 50$, $Y = 1 \text{ N} / \text{mm}^{-2}$, $E = 1000 \text{ N} / \text{mm}^{-2}$, $\nu = 0$) is considered unless stated otherwise.

4.1 Tensile bar: One-dimensional strain-localization

The tapered bar in Figure 2a is clamped at the left-hand side. The bar is subjected to a monotonic tensile loading. The final displacement of 1 mm is applied at the free end in 200 steps with a constant rate of $0.66 \times 10^{-4} \text{ mm} / \text{s}$. Figure 2b shows the results of a mesh refinement study performed with uniform subdivisions of the bar domain along the axis.

Algorithm 7: Online phase

1. Initiate standard FEM code and load:
 - 1.1. the set \mathbf{V} of POD basis vectors associated with the displacement field, Algorithm 6
 - 1.2. the sample matrix of the force vector \mathbf{F}^{samp} and Algorithm 6
the mapping *sampleToClus* between samples and clusters,
 - 1.3. the DEIM coefficient matrices ${}_c\mathbf{A}$, the matrix ${}_c\mathbf{P}$, and Algorithm 6
the DEIM selected entries ${}_c\mathbf{IND}$ for all clusters ($c = 1 \dots n_c$).
2. Set the cluster $c_{\text{new}} = 1$.
3. Break the time domain into steps \rightarrow For each time step t :
 - 3.1. execute the Newton-Raphson scheme \rightarrow For each iteration j :
 - 3.1.0. compute and save the linear part of the stiffness matrix \mathbf{K}^{L} Equation (14)
in the first time step and the first iteration,
 - 3.1.1. compute $\hat{\mathbf{f}} = {}_{c_{\text{new}}}\mathbf{A} \mathbf{f}_{\text{DEIM}({}_{c_{\text{new}}}\mathbf{IND})}$ and $\hat{\mathbf{K}} = {}_{c_{\text{new}}}\mathbf{A} \mathbf{K}_{\text{DEIM}({}_{c_{\text{new}}}\mathbf{IND})}$ Section 3.2
 - 3.1.2. compute the reduced residual vector $\tilde{\mathbf{r}}^j = \mathbf{V}^{\text{T}} (\mathbf{K}^{\text{L}} \mathbf{a}^j + \hat{\mathbf{f}})$ and Section 3.1.1
the reduced stiffness matrix $\tilde{\mathbf{K}}^j = \mathbf{V}^{\text{T}} (\mathbf{K}^{\text{L}} + \hat{\mathbf{K}}) \mathbf{V}$, Section 3.1.1
 - 3.1.3. solve the reduced system of equations $\tilde{\mathbf{K}}^j d\tilde{\mathbf{a}}^{j+1} = -\tilde{\mathbf{r}}^j$, Equation (18.1)
 - 3.1.4. update the reduced solution: $\tilde{\mathbf{a}}^j = \tilde{\mathbf{a}}^j + d\tilde{\mathbf{a}}^{j+1}$, Equation (18.2)
 - 3.1.5. check convergence: $\tilde{\mathbf{r}}^j < \text{tol}$.
 - 3.2. Set the cluster $c_{\text{old}} = c_{\text{new}}$.
 - 3.3. Pass c_{old} , $\hat{\mathbf{f}}$, \mathbf{F}^{samp} , and *sampleToClus* to the nearest neighbor classifier Section 3.3.2
(Algorithm 5) and find the cluster c_{new} that best represents the current state of the problem.
 - 3.4. Go to the next time step.

The mesh refinement study identifies the solution related to the 71 eight-node regular hexahedral element mesh as the reference solution for the development of a reduced order model.




Figure 2: One-dimensional strain-localization problem: (a) the clamped end of the tapered bar has a smaller cross-section to trigger the initiation of plasticity and consequent strain localization; (b) the regularizing properties of viscoplasticity are reflected in the mesh refinement study in terms of the resultant force at the clamped end plotted against the displacement at the free end.

To test the MOR technique, we consider three different values of the b parameter ($b = [0, 20, 100]$). According to (8), large positive values of b result in a fast decay of σ_Y as κ grows and characterize the softening branch of the load-displacement curve as shown in Figure 2b. A reference solution is computed for each value of b , and the corresponding displacement and force samples are collected in matrices $\mathbf{A}^{\text{samp}} = [\mathbf{a}_1 \dots \mathbf{a}_{200}]$ and $\mathbf{F}^{\text{samp}} = [\mathbf{f}_1 \dots \mathbf{f}_{200}]$.

A reduced order model is constructed using these samples for each value of the b parameter. The result of the online computation, depicted in Figure 3, evidently show the complexity reduction achieved by employing our MOR technique. For instance, for the steepest strain softening curve, obtained with $b = 100$, the size of the system of equations is reduced, without apparent loss of accuracy, to 93×93 from 864×864 and, in the light of DEIM, evaluation of the force vector is required to be carried out at 57 entries rather than 864. A subdivision of the samples into three clusters allows to further reduce this number to 63-17-15 (the notation 63-17-15 indicates that we have considered three clusters and that the first cluster requires the evaluation of 63 entries, the second requires 17 entries, and the last only 15). For the other two curves, the size of the system of equations is reduced to 115×115 ($b = 20$) and 94×94 ($b = 0$). The evaluation of the force vector is carried out at 30 and 53 entries, respectively, when employing DEIM. With LDEIM and three clusters, it is possible to further reduce the number of entries to 32-11-8 and 51-22-13, respectively.




Figure 3: The MOR technique applied to the one-dimensional strain-localization problem. The reference curves are obtained with the 71 eight-node regular hexahedral element mesh (864 DOFs) using different values of the b parameter (8).

4.2 Wall under compression: Shear band formation

This example is dedicated to a wall under compression that shows a localization zone in form of a shear band. The lower and the left-hand side edges of the wall depicted in Figure 4a are constrained in the x and y directions. The wall is subjected to a monotonic compressive loading. The final displacement of 2 mm is applied at the upper edge in 400 steps with a constant rate of 1.33×10^{-4} mm / s. The parameter b in (8) is taken equal to 100.

A mesh refinement study has been carried out considering uniform grids with one element across the thickness. Figure 4b shows the results of the mesh refinement study. The solution obtained with 196 (14×14) eight-node regular hexahedral elements is taken as the reference solution.

Displacement and force samples are collected at each time step in matrices $\mathbf{A}^{\text{samp}} = [\mathbf{a}_1 \dots \mathbf{a}_{400}]$ and $\mathbf{F}^{\text{samp}} = [\mathbf{f}_1 \dots \mathbf{f}_{400}]$. Figure 5a shows the convergence of DEIM to the reference solution by a basis addition study. The number of the basis vectors is chosen based on the sum of the square of the neglected singular values as explained in Section 3.1.2. It is possible to make an analogy between adding the basis vectors associated with the singular values shown in Figure 5b for a basis addition study and increasing the number of elements for a mesh refinement study. By adding each subsequent basis vector, the squared value of its associated singular value (say σ_m) is deducted from the error defined in (21) –the error then becomes $\sum_{i=m+1}^r \sigma_i^2$ in which the contribution of σ_m^2 is missing thus yielding a smaller value.

Figure 5b evidently shows the slow rate of decay of the singular values. As a result, a large number of basis vectors is required to decrease the projection error, that is the sum of the square of the neglected singular values, to an acceptable value and render the reduced order model sufficiently accurate. Figure 5a shows that at least 88 DEIM entries (brown curve) are required to construct a reduced order model that is capable of reproducing the reference solution (green curve). The black dashed line is to confirm that we have reached the converged solution. Checking the difference between results of two consecutive basis addition cases can be an adequate criterion to determine the convergence of a reduced order model. It is also noteworthy to report a case where the DEIM approximation results in an unstable reduced order model as shown by the red curve in Figure 5a. The source of this instability is, to the best of our knowledge, still unknown. The same type of behavior is reported in [9] and [13].

The force vector and the stiffness matrix associated with these 88 DOFs are computed by evaluating all the elements shown in Figure 7a, clearly a major subset of the elements in the mesh. As a result, significant reduction in computation cost can not be achieved. To mitigate this issue, we cluster the samples of the force vector that are ‘close’ to each other and compute, for each cluster, a set of reduced basis vectors and a set of DEIM selected entries as discussed in Section 3.3. As shown in Figure 6a, a basis addition study is carried out with 5 clusters. The reduced order model with 19-10-9-9-9 DEIM entries (brown curve) accurately approximates the reference solution (green curve). The singular values corresponding to the x component of the force vector samples are depicted in Figure 6b for each




Figure 4: Two-dimensional strain localization problem: (a) geometry of the wall under compression with the dark shaded part indicating the weak region that localizes the initiation of plasticity and triggers a shear band; (b) the regularizing properties of viscoplasticity are reflected in the mesh refinement study in terms of the resultant of the force in the vertical direction on the lower edge against the value of the displacement of the upper edge; the inset shows the deformed wall with the shear band due to strain localization.




Figure 5: Performance of the DEIM approximation in the two-dimensional strain localization problem: (a) basis addition study (reference solution: 1350 DOFs, 196 elements; reduced order model: 227 POD basis vectors associated with the displacement field); (b) singular values of the force vector samples.

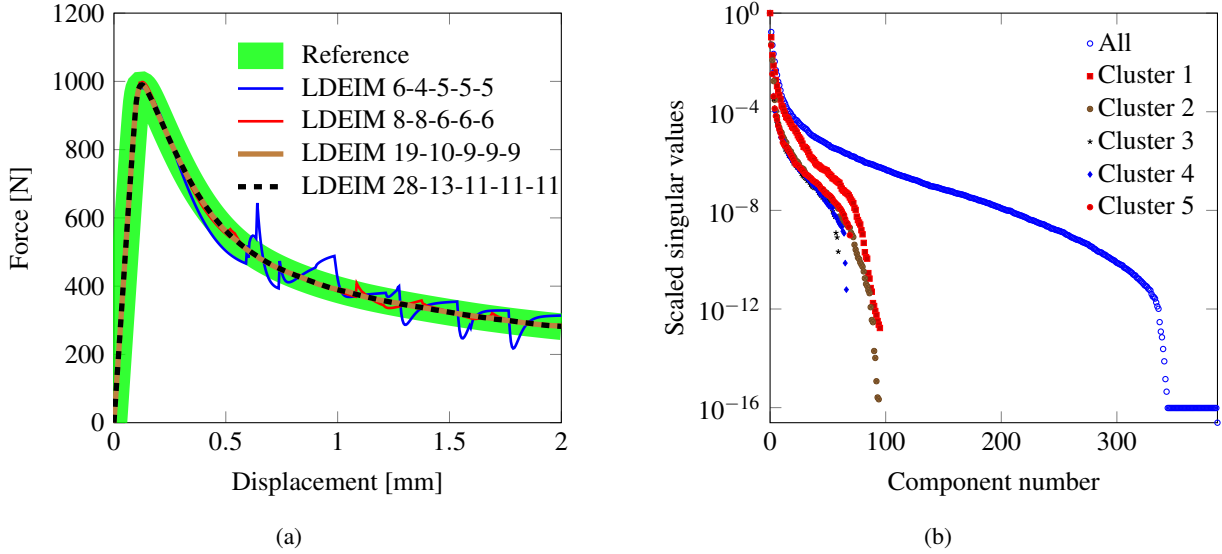


Figure 6: Performance of the LDEIM approximation in the two-dimensional strain localization problem: (a) basis addition study (reference solution: 1350 DOFs, 196 elements; reduced order model: 227 POD basis vectors associated with the displacement field); (b) singular values of the x component of the force vector samples using unclustered and clustered samples (those corresponding to the y and the z components follow the same trend).



Figure 7: Elements used in the online phase in the two-dimensional strain localization problem: (a) DEIM –brown line in Figure 5a; (b) LDEIM –brown line in Figure 6a.

cluster of the reduced order model (those corresponding to the y and the z components follow the same trend). It is clear that the singular values associated with each cluster decay significantly faster than those of all samples.

The elements used in this case are depicted in Figure 7b. Due to the path-dependent nature of the model, all elements depicted in this figure need to be updated for their history parameters. However, only a subset of them is used to compute the force vector and the stiffness matrix at each Newton-Raphson iteration. Figures 7a and 7b show that the online phase of LDEIM is computationally cheaper than that of DEIM.

As discussed in Section 3.3.2, a nearest neighbor classifier is executed at each time step to find the cluster the current force vector belongs to. The classification is carried out based on the $m \ll N$ DEIM-selected entries of the force vector. Due to the non-periodic nature of the mechanical response in this class of problems, clusters do not get repeated during the run time. For instance, if the first cluster is selected at a time step close to the peak of the load-displacement curve, it will not be selected when the load-displacement curve reaches a plateau at the end. Cluster selection being automatic, this observation can be used to monitor the consistency of cluster switching. In other words, repetition of a cluster during the simulation is an indication that either the clustering procedure in the offline phase or the classification procedure in the online phase is inconsistent with the nature of the path-dependent problem. Additionally, this observation can be used to increase the efficiency of the nearest neighbor classifier, Algorithm 5.

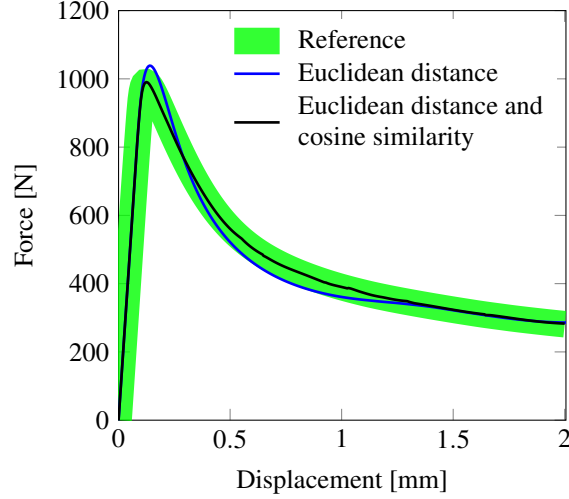


Figure 8: The effect of the distance metric on the load-displacement curves of the two-dimensional strain localization problem (reference solution: 1350 DOFs, 196 element; reduced order model: 227 POD basis vectors associated with the displacement field and LDEIM with 5 clusters and 19-10-9-9-9 DEIM entries in each cluster).

Indeed, by neglecting samples associated with previously selected clusters, the classifier can compare the current state of the system to fewer number of samples.

To achieve consistency as explained in the previous paragraph, we modify the distance metric in the k -means clustering algorithm and the nearest neighbor classifier algorithms. Figure 8 shows that the reduced order model based on the Euclidean distance and the cosine similarity metric (28) (black curve) properly reproduces the reference solution, whereas the reduced order model based on the standard Euclidean distance definition alone fails to do so. For the Euclidean distance case, the cluster number switches from 1 to 5 right before the peak load (note that cluster 5 is associated with the plateau part of the load-displacement curve). This is not consistent and, as a result, the blue curve deviates from the reference equilibrium path.

4.2.1 Parameter sensitivity analysis

We now demonstrate an application of the reduced order model to a simple parameter sensitivity analysis study. The goal is to construct a reduced order model response related to a set of parameters using samples obtained from two different parameter sets. To this end, we run the standard FEM model for $b = 80$ and $b = 120$ (the other parameters are defined at the beginning of Section 4). Samples are uniformly collected in matrices $\mathbf{A}_{b80}^{\text{samp}}$, $\mathbf{F}_{b80}^{\text{samp}}$, $\mathbf{A}_{b120}^{\text{samp}}$ and $\mathbf{F}_{b120}^{\text{samp}}$ at each time step. These samples are combined into two matrices \mathbf{A}^{samp} and \mathbf{F}^{samp} . Since the k -means initial clustering uniformly divides samples into clusters, it is a good practice to combine $\mathbf{F}_{b80}^{\text{samp}}$ and $\mathbf{F}_{b120}^{\text{samp}}$ in a way that samples from the peak load region are collected in the initial columns of \mathbf{F}^{samp} and samples from plateau region are collected in the last columns as described in Algorithm 8.

Algorithm 8: Construct \mathbf{F}^{samp}

input : Two sets of force vector samples for $b = 80$ and $b = 120$: $\mathbf{F}_{b80}^{\text{samp}} \in \mathbb{R}^{N \times n_s}$ and $\mathbf{F}_{b120}^{\text{samp}} \in \mathbb{R}^{N \times n_s}$
output: A mixed sample matrix $\mathbf{F}^{\text{samp}} \in \mathbb{R}^{N \times 2n_s}$
Initiate the mixed sample matrix: $\mathbf{F}^{\text{samp}} = [\mathbf{F}_{b80}(:, 1), \mathbf{F}_{b120}(:, 1)]$
for $i = 2 \dots n_s$ **do**
 Append i th columns of \mathbf{F}_{b80} and \mathbf{F}_{b120} to \mathbf{F}^{samp}
end

A reduced order model is constructed using these samples. In Figure 9a, the reduced order model is employed to simulate the case with $b = 100$. To approximate the standard FEM solution, 85 DEIM entries are required when the

samples are not clustered. However, when the samples are divided into 5 clusters, only 31-19-19-19-17 DEIM entries are required. It is also shown that fewer DEIM entries are needed to reach an acceptably accurate solution (black curve) when the number of clusters increases.

In Figure 9a we observe fluctuations in the load-displacement curves of the reduced order models, starting from a displacement equal to 0.4 mm. A possible explanation could be related to the set of samples that underlies the construction of the basis vectors for the cases with $b = 80$ and $b = 120$. According to (8), the nonlinear parts of the stress related to these two cases, and the corresponding force vectors, differ slightly at the beginning when the plastic strain κ is small; hence, the reduced subspaces corresponding to the two force vectors are close to one another. However, as κ grows, the nonlinear part of the stress for $b = 80$ case decays slower than that for $b = 120$ and, as a result, their corresponding force vectors tend to deviate from one another. Consequently, the corresponding reduced subspaces these two cases belong to are no longer as close as they used to be. At this point, let us consider the case $b = 100$ with a reduced order model constructed using samples collected from the cases with $b = 80$ and $b = 120$. Figure 9a supports the hypothesis just advanced. At the beginning, the load-displacement curves of the reduced order models are fairly close to that of the reference curve (green curve), suggesting that the force vectors for the $b = 100$ case are adequately approximated. Later on, however, small fluctuations in the load-displacement curves of the reduced order models become visible as κ grows.

To reduce these fluctuations, either the number of the basis vectors (and consequently the DEIM entries) in each cluster should be increased or more clusters should be taken into account. The first approach decreases the efficiency of the reduced order model and is not considered. The second approach increases the accuracy of the reduced order model and permits the use of the same or even smaller number of DEIM entries as shown in Figure 9a.

To quantify the impact of the fluctuations in the load-displacement curve of the reduced order model with 10 clusters (black line), we consider the points that coincide with the 8th and 9th cluster change as shown in Figure 9b. These two points correspond to the largest deviation from the reference solution (green curve) and are used to compare their corresponding displacement fields with that of the reference solution employing the relative error measurement

$$e = \frac{|\mathbf{u}_{\text{FEM}} - \mathbf{u}_{\text{LDEIM}}|}{\|\mathbf{u}_{\text{FEM}}\|_2}, \quad (33)$$

where $|\Theta|$ is the component-wise absolute value of Θ . Figure 10 shows the deformed specimen and the relative error corresponding to these two points. It can be observed that the reduced order model is sufficiently accurate in the entire domain with maximum relative error of 3×10^{-4} .

It is noteworthy that we do not enforce clusters to change consecutively. However, we would expect them to change in this manner due to the special initial k -means clustering discussed in Section 3.3.1.

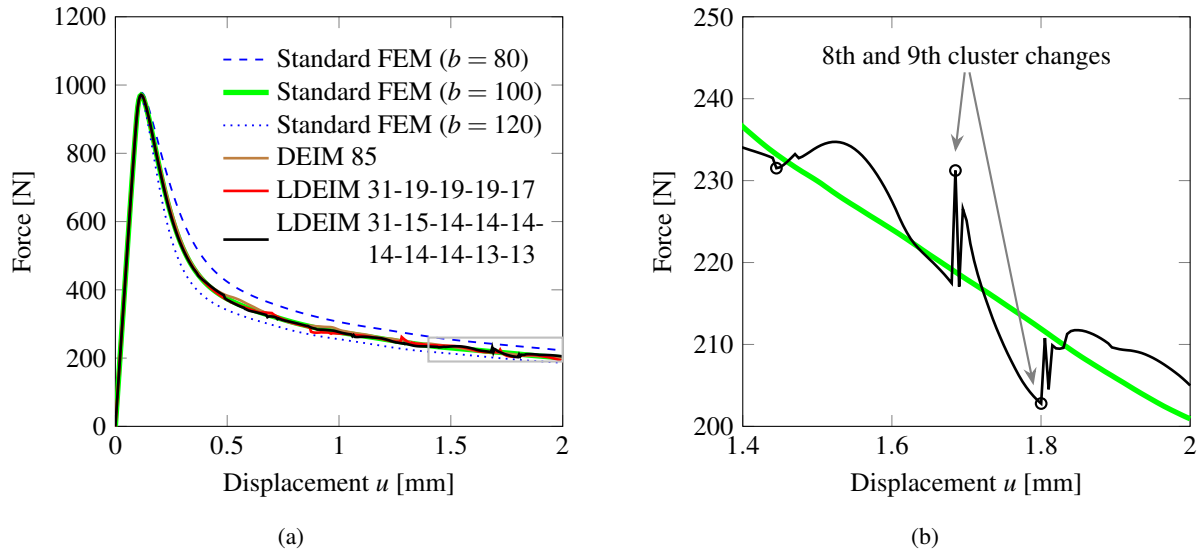


Figure 9: Simple parameter sensitivity analysis for the two-dimensional strain localization problem: (a) evaluation of the case with $b = 100$ using the reduced order model constructed from samples collected from reference cases with $b = 80$ and $b = 120$ (reduced order model: 325 POD basis vectors associated with the displacement field for DEIM and LDEIM); (b) enlargement of the boxed region in (a) with locations of cluster change marked.

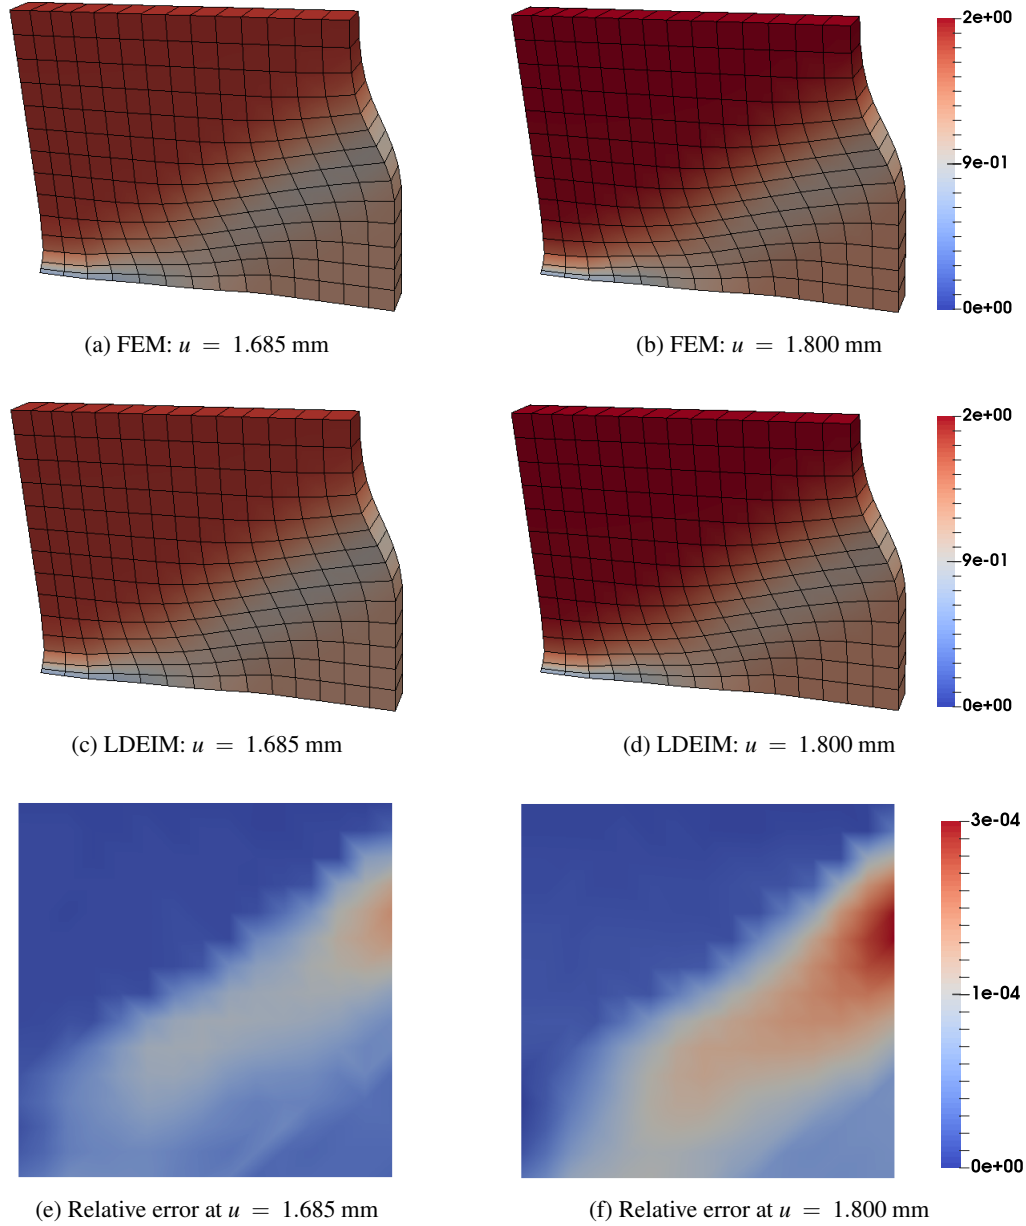


Figure 10: Two-dimensional strain localization problem: deformed configuration corresponding to the black curve depicted in Figure 9a at displacement levels corresponding to the 8th and 9th cluster changes indicated in Figure 9b. The reference FEM solution and reduced order model are depicted in the upper rows and the error between them is depicted in the bottom row.

4.3 Three-dimensional square footing

The footing depicted in Figure 11a is subjected to a monotonic compressive loading. The final compressive displacement of 1 mm is applied at the dark shaded plate on the upper surface in 1200 steps with a constant rate of 0.66×10^{-4} mm / s. The surface on the left-hand side, bottom, and the front are constrained from moving in the x , the z and the y directions, respectively.

A mesh refinement study is carried out and the result is depicted in Figure 11b. Load-displacement curves are associated with the resultant of the forces on the y plane at the bottom and the vertical displacement of the loading plate. Here, we choose the load-displacement curve obtained with 1331 eight-node hexahedral elements as the reference solution. The numerical solution, although not yet converged, is sufficiently accurate for the purpose of this study.

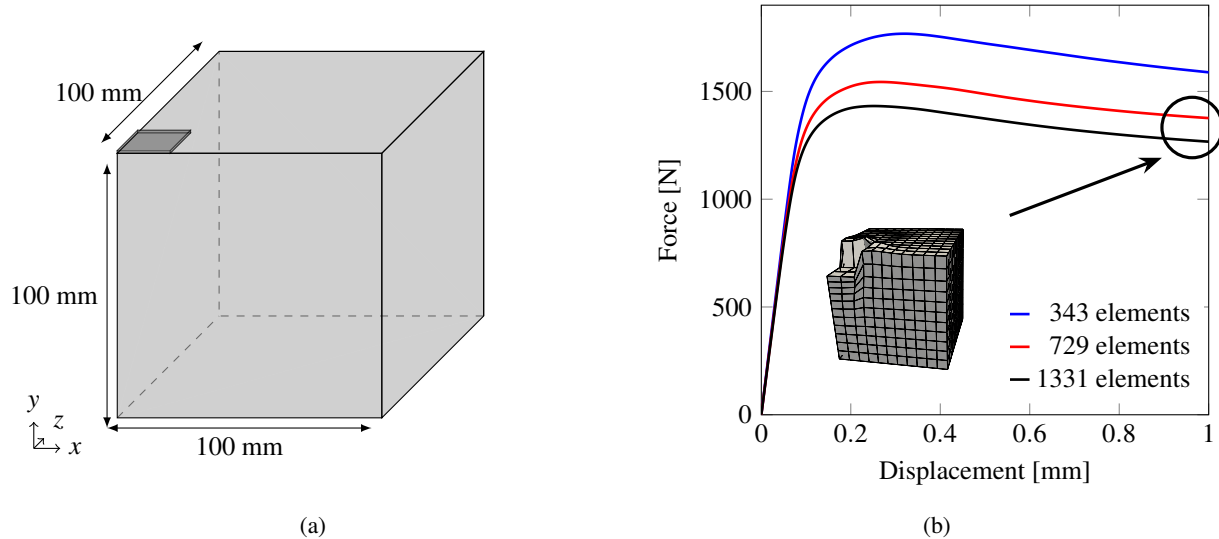


Figure 11: Footing: (a) geometry; (b) mesh refinement study.

In total, 1200 samples were uniformly collected from the reference solution to construct a reduced order model. A large number of POD basis vectors (609), associated with the displacement field, is considered to sufficiently reduce the POD-G projection error. In this manner, it is possible to objectively study the performance of the DEIM approximation.

As shown in Figure 12, at least 335 DEIM entries are required to almost reconstruct the reference solution. Note that the load-displacement curve of the reduced order model starts to deviate from the reference solution before reaching the final displacement value. Nevertheless, it is possible to considerably reduce the number of the DEIM entries by clustering samples. For instance, considering 20, 30, and 40 clusters at most, as depicted in Figure 13, 117, 101, and 43 DEIM entries are required, respectively. This observation implies that by increasing the number of clusters it is possible to construct a significantly more efficient reduced order model. In practice, the amount of memory usage in the online phase is not necessarily proportional to the number of clusters. As depicted in Figure 13, when considering only one cluster, 335 basis vectors are required. Evidently, this number grows to 827 by increasing the number of clusters to 30. However, when the number of clusters becomes 40, the cumulative number of basis vectors shrinks down to 564. The computer memory usage not being necessarily proportional to the number of clusters, considering a large number of clusters is a feasible option.

It is also noteworthy to realize the importance of the distance metric in achieving a consistent clustering. Load-displacement curves for the 40-cluster case (black dashed line in Figure 12) associated with two distance metrics are depicted in Figure 14. It is clear that the metric based on both the Euclidean distance and the cosine similarity definition gives a proper result, whereas the metric based on the Euclidean distance alone fails to do so.

5 Conclusions

Considering the slow decay of the singular values of the samples collected in the offline phase, it is possible to infer that the hypothesis central to all the MOR techniques, reported in Section 1, is not valid for the entire loading history of the class of problems studied in this paper. To mitigate this issue, samples are divided into clusters such that for each cluster the central hypothesis becomes valid. For this purpose, a potentially computationally costly k -means clustering algorithm is added to the offline phase and, employing feature extraction, an efficient nearest neighbor classifier is added to the online phase.

We now identify a possible computational bottleneck pertinent to the reduced stiffness matrix (step 3.1.2 in Algorithm 7) –note that this is merely a technical problem and not a fundamental one as those described at the end of Section 2. The storage of $\mathbf{K} = \mathbf{K}^L + \hat{\mathbf{K}}$ could be quite memory consuming due to the structure of the DEIM approximated nonlinear part $\hat{\mathbf{K}}$. Unlike the linear part of the stiffness matrix \mathbf{K}^L , its sparsity could be considerably less than what is expected in a classical FEM analysis. Additionally, standard procedures that take advantage of the sparsity of

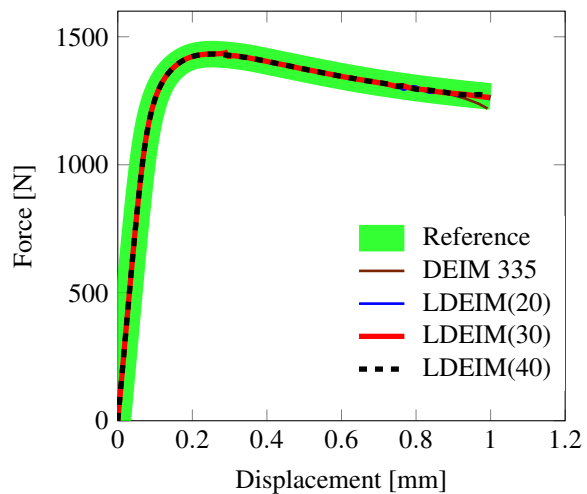


Figure 12: Footing: load-displacement curves associated with different number of clusters (reference solution: 5184 DOFs, 1331 elements; reduced order model: 609 POD basis vectors associated with the displacement field for DEIM and LDEIM).

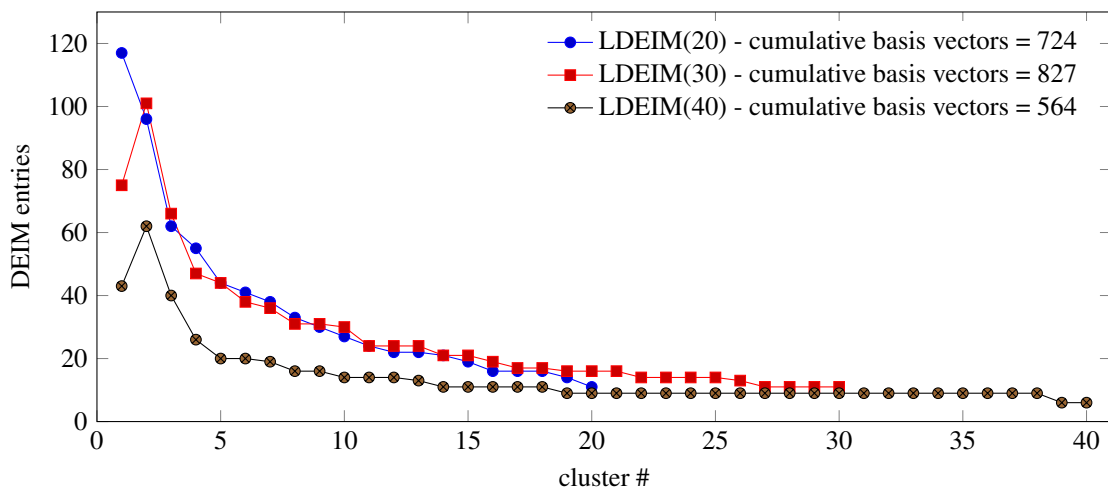


Figure 13: Footing: number of the DEIM entries in each cluster (reduced order model: 609 POD basis vectors associated with the displacement field).

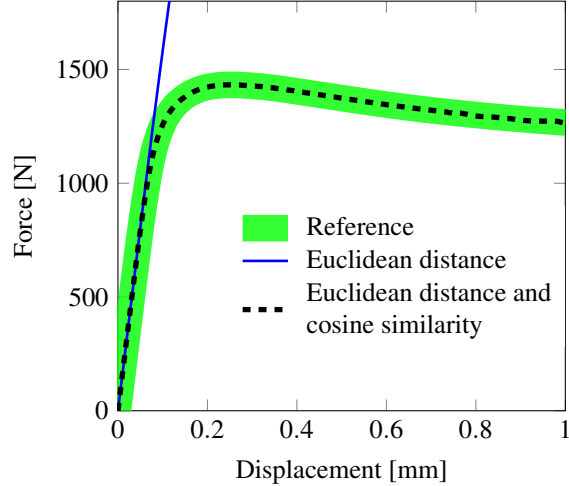


Figure 14: Footing: the effect of the distance definition on the load-displacement diagram (reference solution: 5184 DOFs, 1331 elements; reduced order model: 609 POD basis vectors associated with the displacement field, LDEIM with 40 clusters).

the stiffness matrix, such as matrix assembly, might not be able to perform as efficiently as anticipated in a FEM analysis. As a result, we suggest to avoid explicitly forming $\hat{\mathbf{K}}$ when computing the reduced stiffness matrix by employing an appropriate matrix chain multiplication.

Acknowledgements

The insightful comments of one of the anonymous reviewers is highly appreciated. The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013)/ERC Grant agreement n° 617972.

References

- [1] A. Ammar, B. Mokdad, F. Chinesta, and R. Keunings. A new family of solvers for some classes of multidimensional partial differential equations encountered in kinetic theory modeling of complex fluids. *Journal of Non-Newtonian Fluid Mechanics*, 139(3):153–176, 2005.
- [2] D. Amsallem and B. Haasdonk. PEBL-ROM: Projection-Error Based Local Reduced-Order Models. *Advanced Modeling and Simulation in Engineering Sciences*, 3(1):1–25, 2016.
- [3] A Astolfi. Model reduction by moment matching for linear and nonlinear systems. *IEEE Transactions on Automatic Control*, 55(10):2321–2336, 2010.
- [4] P. Astrid, S. Weiland, K. Willcox, and T. Backx. Missing point estimation in models described by proper orthogonal decomposition. *IEEE Transactions on Automatic Control*, 53(10):2237–2251, 2008.
- [5] B. Besselink, U. Tabak, A. Lutowska, N. van de Wouw, H. Nijmeijer, D. J. Rixen, M. E. Hochstenbach, and W. H. A. Schilders. A comparison of model reduction techniques from structural dynamics, numerical mathematics and systems and control. *Journal of Sound and Vibration*, 332(19):4403–4422, 2013.
- [6] K. Carlberg, C. Farhat, J. Cortial, and D. Amsallem. The GNAT method for nonlinear model reduction: Effective implementation and application to computational fluid dynamics and turbulent flows. *Journal of Computational Physics*, 242:623–647, 2013.

- [7] K. Carlberg, M. Barone, and H. Antil. Galerkin v. least-squares Petrov-Galerkin projection in nonlinear model reduction. *arXiv preprint arXiv:1504.03749v2*, 2015.
- [8] A. Chatterjee. An introduction to the proper orthogonal decomposition. *Current Science*, 78(7):808–817, 2000.
- [9] S. Chaturantabut and D. C. Sorensen. Application of POD and DEIM on dimension reduction of nonlinear miscible viscous fingering in porous media. *Mathematical and Computer Modeling of Dynamical Systems*, 17(4):337–353, 2009.
- [10] S. Chaturantabut and D. C. Sorensen. Nonlinear model reduction via discrete empirical interpolation. *SIAM Journal on Scientific Computing*, 32(5):2737–2764, 2010.
- [11] M. Chevreuil and A. Nouy. Model order reduction based on proper generalized decomposition for the propagation of uncertainties in structural dynamics. *International Journal for Numerical Methods in Engineering*, 89(2):241–268, 2012.
- [12] A. Doostan, R. G. Ghanem, and J. Red-Horse. Stochastic model reduction for chaos representations. *Computer Methods in Applied Mechanics and Engineering*, 196(37-40):3951–3966, 2007.
- [13] C. Farhat, P. Avery, T. Chapman, and J. Cortial. Dimensional reduction of nonlinear finite element dynamic models with finite rotations and energy-based mesh sampling and weighting for computational efficiency. *International Journal for Numerical Methods in Engineering*, 98(9):625–662, 2014.
- [14] S. Gugercin and A. C. Antoulas. A survey of model reduction by balanced truncation and some new results. *International Journal of Control*, 77(8):748–766, 2004.
- [15] B. Haasdonk, M. Dihlmann, and M. Ohlberger. A training set and multiple basis generation approach for parametrized model reduction based on adaptive grids in parameter space. *Mathematical and Computer Modelling of Dynamical Systems*, 17:423–442, 2011.
- [16] S. Herkt, M. Hinze, and R. Pinnau. Convergence analysis of Galerkin POD for linear second order evolution equations. *Electronic Transactions on Numerical Analysis*, 40:321–337, 2013.
- [17] I. T. Jolliffe. *Principal Component Analysis*. Springer, 2nd edition, 2002.
- [18] P. Kerfriden, J. C. Passieux, and S. P. A. Bordas. Local/global model order reduction strategy for the simulation of quasi-brittle fracture. *International Journal for Numerical Methods in Engineering*, 89(2):154–179, 2012.
- [19] P. Krysl, S. Lall, and J. E. Marsden. Dimensional model reduction in non-linear finite element dynamics of solids and structures. *International Journal for Numerical Methods in Engineering*, 51(4):479–504, 2001.
- [20] S. Lall, J. E. Marsden, and S. Glavaški. A subspace approach to balanced truncation for model reduction of nonlinear control systems. *International Journal of Robust and Nonlinear Control*, 12(6):519–535, 2002.
- [21] Q. Liang and L. Wang. Redundancy reduction in wireless sensor networks using SVD-QR. In *Military Communications Conference (MILCOM)*, pages 1857–1861. IEEE, 2005.
- [22] Y. C. Liang, H. P. Lee, S. P. Lim, W. Z. Lin, K. H. Lee, and C. G. Wu. Proper orthogonal decomposition and its applications—Part I: Theory. *Journal of Sound and Vibration*, 252(3):527–544, 2002.
- [23] A. Naik. Data Clustering Algorithms: K-Means Clustering Algorithm, 2010. <https://sites.google.com/site/dataclusteringalgorithms/k-means-clustering-algorithm>. Accessed 15 May 2016.
- [24] B. Peherstorfer, D. Butnaru, K. Willcox, and H. J. Bungartz. Localized discrete empirical interpolation method. *SIAM Journal on Scientific Computing*, 36(1):A168–A192, 2014.
- [25] P. Perzyna. Fundamental Problems in Viscoplasticity. *Advances in Applied Mechanics*, 9:243–377, 1966.
- [26] C. W. Rowley. Model reduction for fluids, using balanced proper orthogonal decomposition. *International Journal of Bifurcation and Chaos*, 15(3):997–1013, 2005.

- [27] D. Ryckelynck. Hyper-reduction of mechanical models involving internal variables. *International Journal for Numerical Methods in Engineering*, 77(1):75–89, 2009.
- [28] P. Tiso and D. J. Rixen. Discrete empirical interpolation method for finite element structural dynamics. In *Topics in Nonlinear Dynamics*, volume 1, pages 203–212. Springer, 2013.
- [29] P. Tiso, R. Dedden, and D. J. Rixen. A modified discrete empirical interpolation method for reducing non-linear structural finite element models. In *ASME 2013 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. American Society of Mechanical Engineers, 2013.
- [30] S. Volkwein. *Proper Orthogonal Decomposition: Theory and Reduced-Order Modeling*. University of Konstanz, 2013. Lecture Notes.
- [31] K. Willcox. Unsteady flow sensing and estimation via the gappy proper orthogonal decomposition. *Computers & Fluids*, 35(2):208–226, 2006.
- [32] K. Willcox and J. Peraire. Balanced Model Reduction via the Proper Orthogonal Decomposition. *AIAA Journal*, 40(11):2323–2330, 2002.