


High-Performance Model Reduction Procedures in Multiscale Simulations

J.A. Hernández
J. Oliver
A. E. Huespe
M. Caicedo



High-Performance Model Reduction Procedures in Multiscale Simulations

**J.A. Hernández
J. Oliver
A.E. Huespe
M. Caicedo**

Monograph CIMNE N°-127, March 2012

INTERNATIONAL CENTER FOR NUMERICAL METHODS IN ENGINEERING
Edificio C1, Campus Norte UPC
Gran Capitán s/n
08034 Barcelona, Spain
www.cimne.com

First edition: March 2012

HIGH-PERFORMANCE MODEL REDUCTION PROCEDURES IN MULTISCALE SIMULATIONS
Monograph CIMNE M127
© The authors

ISBN: 978-84-9939640-6-1

Depósito legal: B-10997-2012

Acknowledgments

The Spanish Ministry of Science and Innovation, and the Catalan Government Research Department, are gratefully acknowledged for their financial support to this research under grants BIA2011-24258 and 2009 SGR 1510, respectively.

Table of Contents

1	Introduction	15
1.1	Motivation	15
1.1.1	Computational multiscale modeling (CMM)	15
1.2	Approach	16
1.2.1	Reduced-order modeling (ROM)	16
1.2.2	High-performance model reduction procedures	18
1.3	Novelty of the approach	20
1.3.1	Main original contribution	20
1.4	Scope	21
2	State of the art	23
2.1	State of the art for multiscale modeling	23
2.1.1	Hierarchical approach	23
2.1.2	Concurrent approaches	25
2.2	State of the art for model reduction	25
2.2.1	Standard projection-based methods	26
2.2.2	Hyper-reduction methods	27
2.3	Combination of multiscale modeling and model reduction	27
3	Standard reduced-order modeling (ROM) of the RVE problem	29
3.1	Homogenization approach	29
3.1.1	Basic assumptions	29
3.1.2	Space of kinematically admissible fluctuations	30
3.1.3	Homogenized stress	31
3.1.4	Incremental variational formulation of the equilibrium of the RVE	31
3.2	Finite element discretization: the full-order model (FOM)	32
3.2.1	Elastic micro-constitutive law	34
3.3	Construction of the reduced-order model (ROM)	34
3.3.1	Sampling of the parametric space	34
3.3.2	Galerkin projection onto the reduced subspace	36
3.3.3	Analysis of approximation errors	37
4	Hyperreduced-order modeling (HROM) of the RVE	41
4.1	Matrix formulation	41
4.2	Approximation of the gauss point stress vector	42
4.3	Ill-posedness of the reduced order problem	43
4.4	Proposed solutions	44

4.4.1	First “tentative” reformulation of the reduced-order model	44
4.4.2	Second tentative reformulation	45
4.4.3	B-matrix orthonormal-based solution; indirect version	46
4.4.4	Definitive solution: expansion of the basis matrix	51
4.5	Analysis of approximation errors	53
4.5.1	Sampling and reconstruction errors	53
4.5.2	Consistency	54
4.6	Computation of macroscopic stresses	55
4.6.1	Homogenized algorithmic tangent moduli	56
5	Assessment of approximation errors: a case study	61
5.1	Introduction	61
5.2	First reduction stage (ROM)	62
5.3	Second reduction stage (HROM)	67
5.3.1	Optimal choice of truncation levels and number of sample points	67
5.3.2	Influence of the greedy sampling strategy	70
5.3.3	Gains in computational efficiency	71
A	The Singular Value Decomposition (SVD)	77
A.1	Introduction	77
A.2	Formulation	77
A.3	Correlation between snapshots: the covariance matrix	78
A.4	Some properties of the SVD	78
B	Influence of mesh refinement on the truncation error	81
B.1	Introduction	81
B.2	Asymptotic analysis	82
B.3	Numerical assessment	82
B.3.1	Discussion of results	83
C	Selection of sample points	87
C.1	Introduction	87
C.2	Selection strategies	88
C.2.1	Statement of the optimization problem	89
C.2.2	Error bound	90
C.2.3	Greedy algorithms	91
C.3	Example	92
C.3.1	Comparison between greedy algorithms	94
	References	96

Nomenclature

α	Vector of N_u random coefficients (solution of the second stage of the model reduction problem)
\mathbf{A}	Finite element assembly operator.
Ψ	Matrix formed by the first N_σ left singular vectors of \mathbf{X}_σ^I .
Ψ_B	Expanded basis matrix
$\hat{\Psi}_B$	“Gappy matrix formed by files appropriately selected from Ψ_B .”
$\hat{\Psi}$	“Gappy matrix formed by selected files of the first N_σ left singular vectors of \mathbf{X}_σ^I .”
Ψ^\perp	Matrix formed by the trailing left singular vectors of \mathbf{X}_σ^I (a basis for the nullspace of \mathbf{X}_σ^{IT}).
Ψ_η	Matrix formed by the first N_σ left singular vectors of \mathbf{X}_σ^I (perturbed stress vector).
$\hat{\Psi}_\eta$	“Gappy matrix formed by selected files of Ψ_η .”
$\partial\Omega_\mu$	Boundary of the domain of the RVE
\mathbf{B}	Matrix connecting strain tensor with nodal displacements.
${}^e\mathbf{B}$	e – th elemental matrix connecting strain tensor with nodal displacements.
\mathbb{B}^{hT}	Matrix that acts on the “stacked gauss point stress vector Σ^I to give the full-order residual vector.
\mathbf{B}^I	Matrix connecting strain tensor with “reduced nodal displacements.
$\hat{\mathbb{B}}^{II^T}$	Matrix that acts on the “gappy gauss” point stress vector $\hat{\Sigma}^{II}$ to give the reduced residual vector.
\mathbb{B}^{IT}	Matrix that acts on the “stacked gauss point stress vector Σ^I to give the reduced residual vector.
$\hat{\mathbb{B}}^{IT}$	Matrix defined as $\hat{\mathbb{P}}\mathbb{B}^{IT}$.
\mathbf{C}_e	Elastic constitutive tensor.
\mathbf{S}_σ^I	Covariance matrix (stress vector snapshots).

\mathcal{C}	Algorithmic tangent operator consistent with the macroscopic incremental constitutive law.
\mathcal{C}_μ	Algorithmic tangent operator consistent with the microscopic incremental constitutive law.
$\mathcal{D}(\bullet)[\Delta u]$	Directional derivative in the direction of Δu .
Ω_μ	Domain of the RVE
Ω_μ	Void part of the RVE domain
Ω_μ^s	<i>Solid</i> part of the RVE domain
${}^e\Omega_\mu^s$	Solid part of the RVE domain corresponding to the e -th element.
ϵ	Macroscopic strain tensor
${}^i\epsilon_n$	Macroscopic strain tensor at t_n (i-th strain history)
ϵ_n	Macroscopic strain tensor at t_n
ϵ_t^{snp}	Macroscopic strain history used for generating the snapshots
$e_\Sigma^{II,rec}$	Error in stresses due to reconstruction (second reduction)
e_Σ^I	Error in stresses (first reduction)
e_Σ^{II}	Error in stresses (attributed exclusively to the second reduction stage)
$e_\Sigma^{I,samp}$	Error in stresses due to deficiencies in the sampling of the parametric space (first reduction)
$e_\Sigma^{II,samp}$	Error in stresses due to deficiencies in the sampling of the parametric space (second reduction)
e_Σ^{trun}	Total truncation error in stresses
$e_\Sigma^{I,trun}$	Truncation error in stresses (first reduction)
$e_\Sigma^{II,trun}$	Truncation error in stresses (second reduction)
$e_{u_\mu}^{I,samp}$	Error in micro-displacement fluctuations due to deficiencies in the sampling of the parametric space (first reduction)
$e_{u_\mu}^{I,samp}{}^\perp$	Component orthogonal to the span of the basis of the sampling error (error in micro-displacement fluctuations due to deficiencies in the sampling of the parametric space (first reduction))
$e_{u_\mu}^{I,trun}$	Truncation error in micro-displacement fluctuations (first reduction)
ϵ_μ	Microscopic strain tensor
η	Perturbation factor (solution of the second stage of the model reduction problem)
\mathbf{f}_ϵ	For the elastic case, $\mathbf{f}_\epsilon \in \mathbb{R}^{N \times d_s}$ is defined as $\mathbf{f}_\epsilon = \int_{\Omega_\mu^s} \{\mathbf{B}^T\} \{\mathbf{C}_\mu\} d\Omega$.

$\hat{\mathcal{I}}$	The “gappy indices set, i.e., the set of indices corresponding to the gauss points selected the <i>gappy data</i> reconstruction.
$\nabla^s A$	Symmetric gradient of A .
\hat{G}_n	Incremental virtual work (or weak form) functional at step n for the micro-equilibrium problem.
I	Identity matrix.
ξ_μ	Microscopic internal variables
\mathbb{A}	Matrix used for expressing the reduced-order model in the form $\mathbb{A}\hat{\Sigma}^{II} = 0$.
t_n^v	Internal traction at time t_n .
\mathbf{K}_μ	Stiffness matrix (micro-structure)
\mathbf{K}_μ^I	<i>Reduced</i> stiffness matrix (micro-structure)
\mathbf{K}_μ^{II}	<i>Hyper-reduced</i> stiffness matrix (micro-structure)
l	Characteristic length of the macroscopic body
l_μ	Characteristic length of the Ω_μ
$\langle \mathbf{A} \rangle$	Mean value of vector A
d	Dimension of the problem ($d = 2, 3$).
N	Number of degrees of freedom of the micro-cell.
n_e	Number of elements of the discretization of the micro-cell.
N_g	Total number of gauss points in the finite element mesh at the coarse level.
N_g	Total number of gauss points in the finite element mesh of the micro-cell.
\hat{N}_g	Number of selected gauss points in the gappy data reconstruction method.
N_h	Total number of strain histories along which displacements and stresses are computed for snapshots construction purposes
N_t	Number of time intervals into which each strain history is divided
N_σ	Number of gauss point vector stress modes (reduced-order model).
N_σ	Number of gauss point vector stress modes (reduced-order model; perturbed snapshot matrix).
N_u	Number of displacement modes.
n	Number of nodes of the discretization of the micro-cell.

$\ A\ _2$	2-norm of A (if A is a vector); if A is a matrix, then $\ A\ _2$ symbolizes the Frobenius norm.)
$\ A\ _\infty$	Infinity norm of A .
N_p	Number of samples (or snapshots) (displacements).
N_s	Number of samples (or snapshots) (stresses).
d_s	Number of components of the strain vector.
null A	Null space or kernel of linear application A
\mathcal{P}_{u_μ}	Tangential operator relating displacement fluctuations and macro-strains.
Φ	Matrix formed by the first N_u left singular vectors of \mathbf{X}_U^h .
${}^e\Phi$	e - th elemental matrix formed by the first N_u left singular vectors of \mathbf{X}_U^h .
$\Phi_{\hat{j}}$	J - th column of Φ .
$\mathbf{R}_{\eta,n+1}^I$	Residual vector at t_{n+1} in the dimensional reduced basis (perturbed problem).
\mathcal{P}_ϵ	For the elastic case, $\mathcal{P}_\epsilon \in \mathbb{R}^{N \times d_s}$ is defined as the operator that acts on ϵ to produce the microscopic fluctuations at interior nodes (for homogeneous boundary conditions).
\mathcal{Q}_{sol}	<i>Linear</i> space that contains the set of all solutions to the equation $\mathbb{B}^{IT} \Sigma^I = \mathbf{T}^{v,I}$.
\mathcal{Q}_{sol}^\perp	Orthogonal complement of \mathcal{Q}_{sol} .
range(A)	Range of linear application A
r_{snp}^u	Rank of the fluctuations snapshot matrix
r_{snp}^σ	Rank of the stress snapshot matrix
rank(A)	Rank of matrix A
\mathcal{R}_σ	Reconstruction matrix (hyper-reduction)
\mathcal{R}_σ	Reconstruction matrix with the expanded basis(hyper-reduction)
\mathbf{R}_I	Residual for node I
\mathbf{R}_{n+1}^{II}	Residual vector at t_{n+1} after the second order reduction.
\mathbf{R}_{n+1}^I	Residual vector at t_{n+1} in the dimensional reduced basis.
\mathbf{R}_{n+1}	Residual vector at t_{n+1}
R_{iI}	i -th component of the residual at node I

$\hat{\mathbb{P}}$	Boolean matrix (containing only zeros and ones) used to extract the files determined by the pertinent greedy algorithm ($\hat{\mathbb{P}} \in \mathbb{R}^{N_{gds} \times \hat{N}_{gds}}$)
N_I	Shape function at node I .
λ_j	j - th singular value computed from the <i>svd</i> of \mathbf{X}_σ^I .
s_μ	Speedup factor associated to the computation of the macroscopic stresses for a given macro-strain history.
s_{macro}	Global speedup factor.
σ_j	j - th singular value (<i>svd</i> of \mathbf{X}_U^h).
Σ	Matrix of singular values (<i>svd</i> of \mathbf{X}_U^h).
σ	Macroscopic stress.
σ_μ	Microscopic stress.
Σ^I	Column vector containing the microscopic stress vector at each gauss point (first reduction).
Σ^h	Column vector containing the microscopic stress vector at each gauss point (full-order model reduction).
Σ_η^I	Column vector containing the ‘‘perturbed microscopic stress vector at each gauss point (first reduction).
Σ_η^I	Column vector containing the perturbation to the microscopic stress vector at each gauss point (first reduction).
Σ^{II}	Column vector containing the microscopic stress vector at each gauss point (second reduction)
$\hat{\Sigma}^{II}$	Column vector containing the microscopic stress vector at the gauss point selected in the gappy data reconstruction.
$\Sigma_{\Psi^\perp}^{II}$	Stress vector normal to the subspace spanned by Ψ .
η	Test functions
η^h	Test functions (finite element approximation)
$\hat{\eta}_I^h$	Nodal test function (at node I)
$\mathbf{T}_n^{v,h}$	Internal traction vector at time t_n (expressed in the final element basis).
$\mathbf{T}_n^{v,I}$	Reduced-order internal traction vector.
\mathcal{T}	Operator that maps micro-stress values at gappy points into the macroscopic stresses
U	Matrix of left singular vectors (<i>svd</i> of \mathbf{X}_U^h).
U_B	Matrix of left singular vectors of \mathbb{B}^I .

\mathbf{U}_Q	Orthogonal basis matrix for \mathcal{Q}_{sol}^\perp .
\mathbf{u}_μ	Fluctuation of the microscopic displacement field
\mathbf{u}_μ^h	Fluctuation of the microscopic displacement field in the finite element approximation
$\hat{\mathbf{u}}_{\mu,i}^h$	Nodal fluctuation displacement at node I .
$\mathbf{u}_\mu _n$	Microscopic displacement at time t_n
\mathbf{u}_μ^I	Displacement fluctuation field pertaining to the reduced space (first reduction)
\mathbf{u}_μ^{II}	Displacement fluctuation field pertaining to the reduced space (second reduction)
$\hat{\mathbf{u}}_\mu^{II}$	Vector of Fourier coefficient corresponding to (displacement fluctuation) basis Φ (second reduction).
$\hat{\mathbf{u}}_\mu^{II} _{\hat{I}}$	Fourier coefficient corresponding to (displacement fluctuation) basis $\Phi_{\hat{I}}$ (second reduction).
$\hat{\mathbf{u}}_\mu^I$	Vector of Fourier coefficient corresponding to (displacement fluctuation) basis Φ (first reduction).
$\hat{\mathbf{u}}_\mu^I _{\hat{I}}$	Fourier coefficient corresponding to (displacement fluctuation) basis $\Phi_{\hat{I}}$ (first reduction).
\mathbf{X}_U^h	Displacement fluctuations snapshot matrix (full-order model).
\mathbf{X}_σ^h	Snapshot matrix of gauss point stress vectors (full-order model).
\mathbf{X}_σ^I	Snapshot matrix of gauss point stress vectors (first reduction).
$\mathbf{U}_{(k)}$	Microscopic nodal displacement vector corresponding to macroscopic strain $(\epsilon)_k$.
\mathbf{u}_μ^m	Microscopic displacement vector at node m
V_μ	<i>Volume</i> of the RVE domain
V_μ^s	<i>Volume</i> of the solid part of the RVE domain
\mathbf{V}	Matrix of right singular vectors (<i>svd</i> of \mathbf{X}_U^h).
\mathcal{K}_μ	Actual set of kinematically admissible microscopic displacements.
$\tilde{\mathcal{K}}_\mu$	Actual set of kinematically admissible displacement fluctuations.
$\tilde{\mathcal{K}}_\mu^*$	Minimally constrained vector space of kinematically admissible displacement fluctuations of the RVE.
\mathcal{K}_μ^I	Minimally constrained vector space of kinematically admissible microscopic displacements.
\mathcal{V}_μ	Vector space of virtual kinematically admissible displacements of the RVE, which happens to be equal to $\tilde{\mathcal{K}}_\mu$.

\mathcal{V}_μ^h	Finite element sub-space (high fidelity solution).
\mathcal{V}_μ^I	Finite element sub-space (first reduced order solution).
w_g	Weight for gauss quadrature at gauss point g
\mathbf{x}	Material point of the macroscopic continuum
\mathbf{y}	Position vector of a point in the microscopic continuum

Chapter 1

Introduction

1.1 Motivation

Technological progress and discovery and mastery of increasingly sophisticated structural materials have been inexorably tied together since the dawn of history. In the present era — the so-called *Space Age* —, the prevailing trend is to design and create new materials, or improved existing ones, by meticulously altering and controlling structural features that span across all types of length scales: the ultimate aim is to achieve *macroscopic properties* (yield strength, ductility, toughness, fatigue limit . . .) tailored to given practical applications. Research efforts in this aspect range in complexity from the creation of structures at the scale of single atoms and molecules — the realm of *nanotechnology* —, to the more mundane, to the average civil and mechanical engineers, development of structural materials by changing the composition, distribution, size and topology of their constituents at the microscopic/mesoscopic level (composite materials and porous metals, for instance).

1.1.1 Computational multiscale modeling (CMM)

The unwavering quest for new, high-performance materials poses, in turn, unprecedented challenges to the 21st century's engineer, who is compelled to devise and develop new modeling tools able to predict, within reasonable accuracy, the mechanical behavior of such materials by taking into account phenomena operating on vastly different scales. Since traditional phenomenological approaches fall clearly short in describing such scale interactions, a new paradigm in computational mechanics has emerged in recent years to address this issue: *computational multiscale modeling* (CMM). Hierarchical methods, sometimes referred to as *sequential or information-passing* methods, are arguably (Fish, 2009) the most widely used CMM techniques. They are based on the *bottom-up*, one-way coupled, description of the material structure (Vernerey et al., 2007). The microstructure of the material is modeled by a sequence of small, nested volumes, each *representative* of a particular scale (the unit cells or representative volume elements (RVE)); starting from the finest level in the microstructure, the response of one scale is computed and infused via cell-averaging techniques (homogenization) into the next coarser level of detail, and so on until arriving at the scale at which engineering predictions are needed (the *macro-*

scopic scale). Multiscale modeling of structural materials using hierarchical homogenization-based strategies has been the subject of intensive research in the past two decades, and formidable advances in the analysis and understanding of scale-bridging mechanism and other multiscale aspects have been reported. However, most of these developments remains still within academic circles.

1.1.1.1 Computational barriers

There is wide consensus (Dolbow et al., 2004) that one of the major factors that hinder exploitation of CMM technology in practical engineering context is the enormous computational cost associated to multiscale simulations. In a two-scale, non-linear history dependent problem discretized using the finite element method (FEM), for example, a microscopic unit cell model (which may contain several thousand of degrees of freedom) has to be solved at every increment and every iteration for each macroscopic (Gauss) point. Furthermore, history data has to be updated at a number of integration points equal to the number of Gauss points in the macro problem multiplied by the number of Gauss points in the unit cell (Fish et al., 1997). Even with the dizzying speed of today’s computers, thus, solution of large structural systems (bridges, nuclear pressure vessels, airplane fuselages, etc...) with accurate resolution of microstructural fields may prove impractical; the situation is aggravated when the system has to be solved many times for various configurations, as occurs in design optimization or inverse analysis, or when more than two scales are involved.

1.2 Approach

1.2.1 Reduced-order modeling (ROM)

This seemingly unsurmountable roadblock — explosive growth of computational cost with the dimension of the problem — is not exclusive of multiscale modeling. Other disciplines dealing with large amounts of multidimensional data, such as data visualization, data mining, image treatment, pattern recognition, etc., also suffer from this computational drawback. The route followed in these disciplines to overcome this barrier is to submit the pertinent data to a process commonly referred to as *dimensionality reduction*¹. Roughly, this process attempts to extract a few *dominant* — in some statistical sense — *structures* or *modes* from a larger data set. In the specific field of computational mechanics, the process of dimensionality reduction is more commonly known as *model reduction*, and it entails the systematic generation of cost-efficient representations of large-scale systems that result, for example, from discretization² via the finite element method (Bui-Thanh

¹In the related literature, the drawback associated to the explosive growth of complexity with the dimension of the problem is (metaphorically) known as the “curse of dimensionality” (Lee and Verleysen, 2007). Extending the metaphor, in CMMS, the curse of dimensionality is induced by the “tyranny of scales” (Oskey and Fish, 2007).

²In the context of mathematical modeling, the terms “discretization” and “reduction” can be considered as conceptually akin, in the sense that both connote transitions from higher to lower solution spaces. Discretization denotes a transition from the continuum description (infinite dimensional space \mathcal{V}) to the discrete approximation (subspace \mathcal{V}^N of

et al., 2007).

The approach followed in this work to attempt to drastically diminish the computational burden associated with multiscale simulations is precisely the *incorporation of model reduction techniques into CMMs*. More specifically, the essence of the proposed approach is to compute the response of the fine scales using *reduced-order equations*; then, following common procedures in hierarchical multiscale modeling, the computed response is infused into the next coarser level of detail via cell-averaging techniques (computational direct homogenization). In this work, the focus is on the solution of the theoretical and algorithmic issues encountered when applying reduced-order modeling techniques to the RVE³ equilibrium problem; other relevant aspects of the multiscale modeling framework are only treated superficially.

1.2.1.1 Reduced-order modeling in multiscale computations

The idea of exploiting the synergistic combination of dimensionality reduction and multiscale modeling is certainly not new. A survey of the related literature reveals that, over the last decade, researchers from various scientific disciplines dealing with multiscale problems have begun to consider in earnest model reduction as a potential route⁴ — complementary to improvements in software and hardware infrastructure — to diminish the often unaffordable cost of multiscale simulations.

In the specific context of homogenization-based multiscale methods, the application of model reduction techniques has been addressed by several authors, namely, Ganapathysubramanian and Zabarar (2004); Yvonnet and He (2007); Boyaval (2007); Monteiro et al. (2008); Nguyen (2008). The strategy adopted in all these works for constructing a cost-efficient model of the micro-cell is the standard *reduced basis method*, whose gist is to project the governing equations onto a low-order subspace spanned by carefully chosen⁵ bases (Amsallem et al., 2009).

1.2.1.2 Limitation of standard reduced-order models

However, reduced basis methods in its standard form suffer from an important limitation when handling non-linear problems: they reduce notably the number of degrees of freedom — and thus the pertinent equation solving effort —, yet the computational cost associated to the evaluation of the inter-

finite dimension N). Model “reduction”, on the other hand, implies that the solution of the equations that govern the pertinent physical system is sought, instead of in \mathcal{V}^N (e.g., the space spanned by the shape functions in the case of discretization via the finite element method), in a subspace \mathcal{V}^M , of dimension M , where $M \ll N$.

³The acronym RVE stands for the “Representative Volume Element”. The goal is to model — in a cost-efficient manner — the evolution of the distribution of stresses within the RVE when subjected to boundary conditions consistent with prescribed macroscopic deformations.

⁴The conclusions of a relatively recent multidisciplinary research initiative sponsored and orchestrated by the USA government attest to this fact: model reduction was pinpointed therein as a strategical research line for the improvement of the overall efficiency of interscale information exchange (Dolbow et al., 2004).

⁵The basis functions are computed in an offline stage from *snapshots* of the solution by means of statistical-based methods such as the Proper-Orthogonal Decomposition (POD).

nal forces and jacobians at quadrature points remains the same⁶. Standard reduction methods only prove, thus, effective in dealing with micro-cells whose constituents are assumed to obey simple constitutive laws (linear elasticity). In a general inelastic case, the calculation of the stresses at each gauss point is, on its own, a computationally expensive operation and dominates the total cost of the computation; as a consequence, the speed up provided by standard model reduction methods in non-linear scenarios is practically negligible, and may not compensate for the cost associated to the *offline* construction of the reduced-order bases.

1.2.2 High-performance model reduction procedures

1.2.2.1 Brief state of the art

Although the standard reduced basis method was introduced in computational mechanics circles in the late 1970s, solutions to overcome the aforementioned limitation were not proposed until relatively recent times. The origin of the first effective proposal can be traced back to the seminal work of Barrault et al. (2004), who suggested to approximate the offending non-linear term in the reduced-order equations by a linear combination of a few, carefully chosen basis functions. In the spirit of the offline/online decomposition of the standard reduced basis approach, these spatial bases are computed *offline* from full-order snapshots of the non-linear term, whereas the corresponding parameter-dependent modal coefficients are determined *online* by *interpolation* at a few (as many as basis functions), judiciously pre-selected spatial points. As in classical reduced bases methods, the efficiency of this *second* or *collateral reduction* is predicated on the existence of a moderate number $M \ll N$ (N is the original dimension of the problem) of basis functions whose span accurately approximate the manifold induced by the parametric dependence of the nonlinear contribution. The interpolation method developed by Barrault et al. (2004) is known as the *Empirical Interpolation Method* (EIM); the main ingredients of this method are: a) the use of a greedy algorithm to generate a set of *maximally independent* bases from the collection of snapshots of the nonlinear term, on the one hand; and b) the recursive selection — also via a greedy algorithm — of spatial locations where the error between the full-order bases and their reconstructed counterparts is maximum.

Other alternative approaches found in the literature are, in essence, similar to that of Barrault et al. (2004) — replacement of the non-linear term by a linear combination of pre-computed basis — and only differ in the manner in which the the basis functions, on the one hand, and the location of the interpolation points, on the other hand, are computed. In the “Best Point” interpolation method (BPIM)⁷, put forward by Nguyen et al. (2008), the bases are generated either on the condition of maximally independence (as in the EIM) or by means of the Proper Orthogonal decomposition (POD); the location of the “best” interpolation points, on the other hand, is deter-

⁶In other words, standard POD/Galerkin methods reduce the *dimension* of the problem, but not necessarily its *complexity* (Chaturantabut and Sorensen, 2010).

⁷We should note that both the EIM and the BPIM were conceived as general-purpose interpolation schemes for parametrized, multidimensional functions.

mined by minimizing the error⁸ between the full-order snapshots and the reduced-order snapshots — the snapshots “reconstructed” from the values at the selected points. As opposed to the EIM, which employs an heuristic, sub-optimal greedy algorithm, in the BPIM, this optimization problem is attacked using an “exact algorithm” (Levenberg-Marquardt scheme); as a consequence, the approximation obtained with the BPIM is, in general, more accurate than the one provided by the EIM⁹, although, obviously, at a considerably higher offline cost. The proposal by Astrid¹⁰ and co-workers (Astrid, 2004; Astrid et al., 2008) for estimating the modal coefficients in the approximation of the nonlinear term is the so-called *Missing Point Estimation* (MPE) method. This method uses POD bases and, as in the EIM, a greedy algorithm to determine the optimum location of the sample points. However, unlike the EIM, the optimization is carried out by minimizing the condition number of the gram matrix¹¹ of the POD bases evaluated at the candidate points. Another notable difference with respect to the EIM (and also the BPIM) is that the MPE can handle situations in which the number of candidate points is greater than the number of bases in the expansion. In such cases, the system of linear equations that permits to obtain the modal coefficients is overdetermined, and recourse to a least-square solution is to be made to resolve the indeterminacy¹². This strategy — least-square regression combined with POD bases — was originally proposed by Everson and Sirovich (1995) for purposes of recovering full images from images marred by *gaps* in the data, and it is known in the related literature as the *gappy* POD method.

In solution methods in which the pertinent governing equations are used in its variational form (as in FE), reduction in complexity arising from nonlinearities can be alternatively achieved by approximating the *integrals* in which the offending nonlinear function appears — rather than the function itself, as done in the interpolatory and least-square reconstruction techniques discussed above. Based on this observation, An et al. (2009) propose a quadrature scheme devised for fast-run integration of the subspace

⁸More precisely, the minimization of the projection of this error onto the space spanned by the selected bases.

⁹See Nguyen et al. (2008); Galbally et al. (2010) for thorough comparisons between both methods.

¹⁰Note that Astrid’s first publication on this topic dates back to 2004 (her doctoral thesis), which is the same year in which Barrault’s seminal work on interpolation applied to model reduction (Barrault et al., 2004) was issued; furthermore, no references to Barrault’s work appears in Astrid (2004) (and vice versa). It seems, thus, that both approaches were conceived and developed independently.

¹¹Let Ψ the POD basis matrix and $\hat{\Psi}$ the sub-block matrix of Ψ containing only the rows associated to the candidate points. The gram matrix or gramian of $\hat{\Psi}$ is defined as $\hat{M} = \hat{\Psi}^T \hat{\Psi}$. Certain choices of candidate points may render this matrix non-invertible, and hence, an effective route to avoid this is to drive a greedy algorithm that minimizes $\text{cond}(\hat{M})$ (Willcox, 2006).

¹²Let $\Sigma(\mathbf{x}, \epsilon) = \Psi(\mathbf{x})\mathbf{c}(\epsilon)$ be the linear expansion of the non-linear, parametrized (in terms of ϵ) function Σ . In an interpolatory scheme, the Fourier coefficients \mathbf{c} are simply calculated by solving the square system $\hat{\Psi}\mathbf{c}(\epsilon) = \hat{\Sigma}(\epsilon)$, where $\hat{\Psi}$ and $\hat{\Sigma}$ denotes the basis matrix and the non-linear function, respectively, evaluated at the selected points. If the number of points is greater than the number of bases, the system $\hat{\Psi}\mathbf{c}(\epsilon) = \hat{\Sigma}(\epsilon)$ becomes overdetermined and only a solution in the least-square sense can be provided: find $\mathbf{c}(\epsilon)$ minimizing $\|\hat{\Psi}\mathbf{c}(\epsilon) - \hat{\Sigma}(\epsilon)\|^2$ (Quarteroni et al., 2000).

spanned by a representative set of snapshots of the nonlinear integrand¹³. In this scheme, the choice of the quadrature points is guided by a greedy algorithm that minimizes the integration error over a set of representative training samples, whereas the corresponding quadrature are calculated using a nonnegative least squares (NNLS) algorithm.

In what follows, we shall use interchangeably the appellations *High-Performance Reduced Order Modeling* (HP-ROM) and Hyper-Reduced¹⁴ Order Modeling (HROM) to refer to reduced basis methods combined with interpolatory or least-square reconstruction schemes.

1.3 Novelty of the approach

In several areas of mathematical modeling and computational mechanics, recognition of the tremendous gain in computational efficiency afforded by these enhanced model reduction techniques has instigated an intense flurry of activity in the advancement and use of such techniques: general parametrized partial differential equations (Grepel et al., 2007; Nguyen, 2007; Rozza et al., 2007; Huynh and Patera, 2007; Nguyen and Peraire, 2008); neuronal modeling (Chaturantabut and Sorensen, 2010); nonlinear viscous fingering in miscible fluids (Chaturantabut and Sorensen, 2010); structural dynamics optimization procedures (Amsallem et al., 2009); uncertainty quantification in inverse problems (Galbally et al., 2010); nonlinear transient dynamic analysis (Carlberg et al., 2011), to name a few. However, to the best of the authors' knowledge, this research activity has not yet penetrated into the specific field of computational multiscale modeling. The present work is thereby intended to take a first step in this direction by exploring and addressing the challenges that may arise in applying these nascent model reduction techniques to multiscale modeling.

1.3.1 Main original contribution

The theory underlying high-performance model reduction — i.e., model reduction using projection methods combined with interpolatory/least-square reconstruction schemes — is still at its embryonic stage of development, and many fundamental issues remains to be addressed. Foremost among these is the question of the well-posedness of the reduced-order problem: does the expansion of the offending non-linear term leads invariably to a well-posed problem? Or is there, on the contrary, any situation in which this desirable characteristic is not present? Examination of the scarce literature on the subject indicates that no researcher has so far been confronted with ill-posed reduced-order equation; in view of this fact, one may be certainly inclined to believe that uniqueness of solution can be taken for granted whenever the full-order model is well-posed. Unfortunately, this is not so, and a clear example in which this proviso is not met is encountered, in the multiscale

¹³As noted by An et al. (2009), this is inspired in Gauss-Legendre quadrature schemes, which are constructed so that functions residing in polynomial spaces can be integrated exactly if values of the integrand are known at an appropriate number of points — more precisely, polynomials up to order $2N - 1$ are integrated exactly using *only* values of the integrand at N strategically selected points.

¹⁴This term was coined by D. Ryckelynck (Ryckelynck, 2005), another fertile contributor to this approach.

context, when attempting to reduce the complexity of the micro-cell equilibrium¹⁵ equations.

Indeed, the variational equation that governs the equilibrium of the micro-cell reads as:

$$\int_{\Omega_\mu^s} \boldsymbol{\sigma}_\mu(\mathbf{u}_\mu(\mathbf{x}), \boldsymbol{\epsilon}) : \nabla^s \boldsymbol{\eta}(\mathbf{x}) d\Omega = 0, \quad \forall \boldsymbol{\eta} \in \mathcal{V}_\mu, \quad (1.3.1)$$

where¹⁶ \mathbf{u}_μ stands for the displacement fluctuation (the state variable), $\boldsymbol{\epsilon}$ denotes the *prescribed* macro-strain tensor (the parameter) and $\boldsymbol{\sigma}_\mu$ represents the stress tensor (the non-linear function). Approximating $\boldsymbol{\sigma}_\mu$ as a linear combination of M pre-computed spatial bases:

$$\boldsymbol{\sigma}_\mu(\mathbf{u}_\mu(\mathbf{x}); \boldsymbol{\epsilon}) \approx \sum_{i=1}^M \boldsymbol{\Psi}_i(\mathbf{x}) c_i(\boldsymbol{\epsilon}, \mathbf{u}_\mu), \quad (1.3.2)$$

and inserting this approximation into the equilibrium equation (1.3.1), we arrive at:

$$\begin{aligned} & \int_{\Omega_\mu^s} \sum_{i=1}^M \boldsymbol{\Psi}_i(\mathbf{x}) c_i(\boldsymbol{\epsilon}, \mathbf{u}_\mu) : \nabla^s \boldsymbol{\eta}(\mathbf{x}) d\Omega \\ &= \sum_{i=1}^M \left(\int_{\Omega_\mu^s} \boldsymbol{\Psi}_i(\mathbf{x}) : \nabla^s \boldsymbol{\eta}(\mathbf{x}) \right) c_i(\boldsymbol{\epsilon}, \mathbf{u}_\mu) = 0, \quad \forall \boldsymbol{\eta} \in \mathcal{V}_\mu. \end{aligned} \quad (1.3.3)$$

By construction, each basis $\boldsymbol{\Psi}_i$ is a linear combination of stress fields $\{\boldsymbol{\sigma}_\mu(\mathbf{x}, \boldsymbol{\epsilon}_j)\}$, ($j = 1, 2, \dots$) that fulfill the equilibrium equations (the snapshots); thus, the bases themselves are (self-)equilibrated fields:

$$\int_{\Omega_\mu^s} \boldsymbol{\Psi}_i(\mathbf{x}) : \nabla^s \boldsymbol{\eta}(\mathbf{x}) = 0, \quad \forall \boldsymbol{\eta} \in \mathcal{V}_\mu. \quad (1.3.4)$$

In view of this equation, it becomes apparent that the equilibrium problem in its reduced-order format (1.3.3) is patently *ill-posed*, since the equilibrium condition is satisfied regardless of the value of the modal coefficients c_i , and thus, regardless of the value of the micro-displacement fluctuations \mathbf{u}_μ . The *main original contribution of the present work to the field of model reduction is a refined coefficient-approximation strategy that safely avoid this type of ill-posedness.*

1.4 Scope

Multiscale modeling is an exceptionally broad field in science, and attempting to devise an all-embracing model reduction strategy would be too ambitious a goal. For this reason, attention is confined in this work on the study of the dimensionality and complexity reduction of models based on one of the most popular — at least within the field of structural materials — class of multiscale hierarchical approaches: the *direct computational homogenization*. Nevertheless, our deliberations may be of aid in illuminating the path

¹⁵In the absence of internal tractions.

¹⁶ \mathcal{V}_μ is the vector space of kinematically admissible RVE displacement fluctuations.

towards the development of model reduction methods for other multiscale approaches; likewise, the methodology proposed here to eliminate the aforementioned type of ill-posedness can be applied to solve similar problems arising in other modeling contexts.

Chapter 2

State of the art

The purpose of this work is to exploit the benefits that may accrue from the merge of two distinct, but complementary, lines of research: computational multiscale modeling (CMM) and reduced-order modeling (ROM). Accordingly, the state-of-the-art review presented in the sequel will be structured in three sections. The first part provides an overview of the current state of the art in multiscale modeling, placing special emphasis on the direct computational approach, as is the one addressed in this work. The second part describes recent advances in the field of model reduction; and, finally, the last section outlines the few attempts made to date in the computational mechanics community to bring the nascent idea of combining CMM and ROM to fruition.

2.1 State of the art for multiscale modeling

Multiscale modeling approaches can be broadly classified (Fish, 2009) as either *hierarchical or concurrent*, according to how information from fine scale models is incorporated into coarse scale models.

2.1.1 Hierarchical approach

Hierarchical methods, sometimes referred to as *sequential or information-passing* methods, are based on the *bottom-up*, one-way coupled, description of the material structure (Vernerey et al., 2007). The microstructure of the material is modeled by a sequence of small, nested volumes, each representative of a particular scale; starting from the finest level in the microstructure, the response of one scale is computed and infused via cell-averaging techniques into the next coarser level of detail, and so on until arriving at the scale of practical interest (the *macroscopic* scale). Two strands of works on hierarchical methods occupy at present the limelight in multiscale modeling: direct *computational homogenization* approaches, and direct *mathematical homogenization* approaches (Yuan and Fish, 2009).

2.1.1.1 Direct computational homogenization

The theoretical foundations of the *direct computational homogenization approach* were laid in the work of Germain (1973). The basic building blocks of such an approach is : 1) the Hill-Mandel Principle of Macro-Homogeneity

(Hill, 1963), whereby the stress power at a point of the coarser scale must equal the volume average of the stress power at the finer scale over a representative volume element (RVE) of the material associated with that point; 2) the assumption that the strain at a point of the coarser scale are the volume average of the fine-level strain field at the RVE. Methods based on this approach have soared in popularity over recent years due to its suitability for implementation in finite element codes (see for instance Michel et al. (1999); Nemat-Nasser (1999); Miehe et al. (2002); Terada et al. (2003); Miehe et al. (1999, 2002); Terada et al. (2003); Giusti et al. (2009)). A general variational formulation for both small and large strain multi-scale solid constitutive models based on this approach is thoroughly described in de Souza Neto and Feijóo (2006); the distinguishing feature of this work is that it provides a clearly articulated axiomatic framework in which the notion of *kinematically admissible RVE displacement fluctuation field* — another key ingredient of direct computational homogenization methods — is formalized.

The Representative Volume Element and the Unit Cell

As pointed out above, at the heart of the direct computational homogenization lies the notion of representative volume element (RVE). The RVE concept is closely related to the so-called *unit cell*, employed in other multi-scale approaches (direct mathematical homogenization). According¹ to Tu et al. (2009), “the unit cell contains a certain physical volume of microstructure, from which continuum quantities (the critical parameters) are computed. The representative volume element (RVE), on the other hand, is defined as the smallest possible region representative of the whole heterogeneous media, on average. Unlike the RVE, the unit cell may not necessarily represent the behavior of the entire domain. However, similar to the RVE, the unit cell is a finite physical domain where a continuum description is applicable.”

The analysis of the effect of heterogeneities in metals is a paradigmatic case that calls naturally for multi-scale analysis Efendiev and Hou (2009). At its simplest, the “heterogeneity” in a macroscopic domain of a given material is the presence of a single void; the renowned Gurson yield criterion (Gurson et al., 1977) for porous materials was derived by a semi-analytical method using an unit cell consisting of a Von Mises matrix with a single void. It should be stressed that an aspect of paramount importance in the construction of multi-scale models using the RVE concept is related with the *boundary conditions* to be imposed on the RVE; de Souza Neto and Feijóo (2006) and Giusti et al. (2009) provides a detailed treatment of this key issue.

2.1.1.2 Direct mathematical homogenization

The *direct mathematical homogenization approaches* are based on the so-called *mathematical homogenization theory*, which has its point of departure in the works by Sánchez-Palencia (Sanchez-Palencia, 1974, 1980). These

¹See also Gross and Seelig (2011), page 245.

works pioneered a new direction, parallel to the theoretical path paved by Germain (1973), in the understanding of how a homogenized medium can be substituted by a non-homogeneous material. In the mathematical homogenization theory, the homogenized equations are the limit equations obtained as a certain dimensionless parameter — the ratio between the characteristic length of the fine scale to the characteristic length of the coarse scale — goes to zero (Michel et al., 1999). Among the several strategies that have been specifically devised to evaluate this limit, the most widely developed are the *multiscale asymptotic expansion methods*, in particular asymptotic expansion methods for periodic heterogeneous structures. A comprehensive survey of these methods is contained in Fish (2009) (chapter 7). One of the most prolific contributors to the field of asymptotic expansion-based homogenization is the research group of J. Fish and co-workers (see for instance Fish and Shek (1999); Fish et al. (2002); Fish and Chen (2004); Chen and Fish (2006); Fish et al. (2007); Yuan and Fish (2008)). A line of approach which has been actively pursued by Fish’s group in recent years is concerned with the *mathematical homogenization with eigenstrains based on the Transformation Field Analysis* (Fish et al., 1997; Oskay and Fish, 2007; Yuan and Fish, 2009; Fish, 2009). The Transformation Field Analysis (Dvorak et al., 1994) allows precomputing certain information in a “preprocessing” phase and hence affords considerable reduction of the computational cost of the direct homogenization approach. The gist of this approach, thus, is somewhat the same as the one pursued in this work.

2.1.2 Concurrent approaches

Multiscale methods in which the fine- and coarse-scale models are executed in concurrent, two-way coupled fashion — as opposed to hierarchical methods, which are sequential, one way-coupled — are called *multiscale concurrent methods*. In these methods, the coarse-scale model serves primarily as a *boundary model* (Xiao and Belytschko, 2004) for the fine-scale problem; compatibility and momentum balance are enforced across the interface between the fine and coarse scales, which, in turn, should be properly contrived to faithfully model information transfer between domains at minimal computational expense. Furthermore, to be computationally tractable, the subdomain where a higher order theory is required (fine-scale model) should be considerably small compared to the domain of the problem. A comprehensive review on multiscale concurrent methods can be consulted in chapter 4 of Fish (2009).

2.2 State of the art for model reduction

The problem of building simpler models, i.e., models comprising fewer and computationally faster to solve equations, from complex ones is generically known as *model reduction*. Using physical insight and intuition, for instance, an initially complex model can be transformed into a more parsimonious one by stripping away the “inessentials” and retaining only the relevant — for practical purposes — physics of the problem; most of the classical Strength of Materials equations were inferred from Elasticity Theory following this approach. At the opposite extreme of this physical insight-based approach

are the so-called *black-box methods* (Astrid, 2004), in which correlations between a comprehensive range of input data and the corresponding output information computed with the complex model are automatically constructed using non-linear statistical data modeling tools such as artificial neural networks (see for instance Lopez et al. (2008); Lopez (2008)). In contrast to the physical insight-based approach, black-box models can be derived in a generic and systematic manner, not contingent upon the skill and physical intuition of the modeler. However, black-box models are derived without taking into account explicitly the structure of the governing equations (in fact, they are more suited for cases in which such equations are not available); thus, pronounced departures from the values of input data used to “train” the model may adversely affect the predictive capacity of black-box models.

2.2.1 Standard projection-based methods

A class of model reduction techniques that have gained prominence in recent years, and that somewhat combine advantageous features of both physical-insight and black box approaches, is the commonly referred to as *projection-* (or reduced basis)- *based methods*. Projection-based methods employ a previously computed set of state solutions (snapshots) to generate a relatively low-dimensional basis whose corresponding subspace intends to approximate, in a certain sense, the full-order solution space (*offline* stage); then, the governing equations are projected onto this reduced-order subspace, resulting in a model with a significantly reduced number of degrees of freedom — the *reduced-order model* (ROM) — that is solved in the *online* stage.

The method that combines the *proper orthogonal decomposition* (POD) (also known as the Karhunen-Loeve expansion, method of empirical eigenfunctions or principal component analysis) and the Galerkin projection is, arguably, the most popular model reduction technique in the computational mechanics community — see Krysl et al. (2001) for a rigorous and extensive discussion on this method and its applicability in non-linear finite element models of solid mechanics. Its popularity can be chiefly attributed to the *optimal properties*, in the least-square sense, of the reduced basis obtained with the POD (Chatterjee, 2000). Variants of the POD/Galerkin model reduction strategy that have gained also certain status in computational mechanics are the *balanced truncation* (Cortelezzi and Speyer, 1998) and the *balanced POD* techniques (Rowley, 2005) scheme.

Due to the obvious computational burden associated to the offline stage (computation of the snapshots and POD), the partitioned offline-online strategy on which these model reduction methods are based proves advantageous only in one of the two following contexts (Carlberg and Farhat, 2008): 1) applications that require “real-time” predictions (interactive design, control); 2) in “many-query” situations, i.e., when the system is solved many times for various configurations. Multiscale modeling, for instance, clearly falls within this latter category, for it requires the repetitive solution of the unit cell problem under varying conditions.

2.2.2 Hyper-reduction methods

Unfortunately, when a general nonlinearity is present, it can be shown (Chaturantabut and Sorensen, 2010) that the computational cost of evaluating the coefficients appearing in the equations of the reduced model still depends on the dimension of the original system (N). Stated alternatively, the standard POD/Galerkin method affects the *dimensionality or order* of the model, but not its *computational complexity*, which still depends on N . Several refinements of standard POD/Galerkin model reduction have been proposed in recent years to overcome this deficiency. The starting point in both approaches is the reduced-order model generated by the POD/Galerkin method, whose nonlinear terms are subjected to a *further* reduction in complexity (hence the appellation *hyper-reduction* methods, coined by Ryckelynck (2005)). A review of the literature on hyper-reduction methods was already presented in section 1.2.2.1.

2.3 Combination of multiscale modeling and model reduction

The literature review on application of model reduction strategies in direct *computational* homogenization methods was already sketched in section 1.2.1.1; we saw that all research efforts to date have been limited to exploit the classical POD/Galerkin model reduction method — with no additional reduction in complexity — to treat the unit cell problem. A distinct line of approach — not related with any of the general projection-based model reductions methods described previously — is the method proposed by J. Fish and co-workers Oskay and Fish (2007); Yuan and Fish (2009); Fish (2009). They propose to apply the so-called Transformation Field Analysis (TFA) to the asymptotic expansion-based homogenization² methods. The TFA is a strategy that allows precomputing certain information (localization operators, concentration tensors, transformation influence functions) in a *preprocessing* phase prior to nonlinear analysis, which consequently, can be carried out with a small subset of unknowns.

²Fish's approach, thus, can be only apply to direct mathematical homogenization problems.

Chapter 3

Standard reduced-order modeling (ROM) of the RVE problem

3.1 Homogenization approach

As pointed out earlier, two lines of approach dominate at present the scenery of multiscale hierarchical modeling: the direct *computational homogenization* approach, and the direct *mathematical homogenization* approach. The model reduction techniques developed in the present work will be applied to multiscale models derived from the former approach —direct *computational* homogenization. In the following, the fundamental assumptions upon which this approach, in its more conventional form, rests are described, although without dwelling unduly on formal considerations; for a more in-depth description of the underlying axiomatic framework, readers are referred to de Souza Neto and Feijóo (2006).

3.1.1 Basic assumptions

1. Existence of a *RVE*.

Associated to every point \mathbf{x} of the macro-continuum $\mathcal{B} \subset \mathbb{R}^d$ ($d = 2, 3$), of characteristic length l , there is a microstructural domain $\Omega_\mu \subset \mathbb{R}^d$ of characteristic length $l_\mu \ll l$ that is *representative* of the heterogeneous material as a whole —the so-called *Representative Volume Element* (RVE). In general, the domain of the RVE is assumed to consist of a solid part Ω_μ^s and a void¹ part Ω_μ^v ($\Omega_\mu = \Omega_\mu^s \cup \Omega_\mu^v$).

2. Additive decomposition of the microscopic strain tensor (*first-order* homogenization).

The microscopic strain tensor $\boldsymbol{\epsilon}_\mu$ at each point \mathbf{y} of the RVE is expressible (presuming infinitesimal deformations) as:

$$\boldsymbol{\epsilon}_\mu(\mathbf{y}) = \boldsymbol{\epsilon} + \nabla^s \mathbf{u}_\mu(\mathbf{y}), \quad (3.1.1)$$

¹In this work, we shall further consider only RVEs whose void part does not intersect the external RVE boundary.

where $\boldsymbol{\epsilon}$ stands for the *macroscopic* strain tensor and \mathbf{u}_μ denotes the microscopic displacement fluctuation.

3. Homogenized strain tensor.

The macroscopic strain tensor $\boldsymbol{\epsilon}$ is assumed to be the *volume average* over the RVE of the microscopic strain field $\boldsymbol{\epsilon}_\mu = \boldsymbol{\epsilon}_\mu(\mathbf{y})$:

$$\boldsymbol{\epsilon} = \frac{1}{V_\mu} \int_{\Omega_\mu} \boldsymbol{\epsilon}_\mu \, d\Omega, \quad (3.1.2)$$

where V_μ is the volume of the RVE:

$$V_\mu = \int_{\Omega_\mu} d\Omega. \quad (3.1.3)$$

4. Hill-Mandel principle of macro-homogeneity.

Let $\boldsymbol{\sigma}_\mu$ be a microscopic stress field in equilibrium. The following identity must hold

$$\boldsymbol{\sigma} : \delta\boldsymbol{\epsilon} = \frac{1}{V_\mu} \int_{\Omega_\mu} \boldsymbol{\sigma}_\mu : \delta\boldsymbol{\epsilon}_\mu \, d\Omega. \quad (3.1.4)$$

for any virtual (i.e., kinematically admissible) microscopic strain $\delta\boldsymbol{\epsilon}_\mu$.

3.1.2 Space of kinematically admissible fluctuations

Equation (3.1.2) can be alternatively expressed in terms of *boundary displacement fluctuations*. Indeed, inserting Eq.(3.1.1) into Eq.(3.1.2), and using the Gauss theorem, one gets:

$$\int_{\partial\Omega_\mu} \mathbf{u}_\mu \otimes_s \mathbf{n} \, d\Omega = \mathbf{0}. \quad (3.1.5)$$

Here, $\partial\Omega_\mu$ stands for the boundary of Ω_μ and \mathbf{n} denotes the outer unit normal vector to $\partial\Omega_\mu$. Note that the set of all displacement fluctuation fields satisfying this condition is a *vector space*; this peculiarity will greatly facilitates the application of model reduction techniques to the RVE equilibrium problem. The different classes of multiscale models discussed in the related literature arise from choosing a particular subspace—henceforth denoted by \mathcal{V}_μ —of this larger space of *kinematically admissible* displacement fluctuations. The most trivial case—but also the most inaccurate—is the renowned Taylor’s model (or “Rule of Mixtures”), which assumes that displacement fluctuations are zero for all $\mathbf{y} \in \Omega_\mu$ ($\mathcal{V}_\mu \equiv \{\mathbf{0}\}$). More reliable (and less kinematically restricted) models only prescribe the displacement fluctuations on the boundary of the RVE, and, consequently, as opposed to Taylor’s model, the equilibrium equation of the RVE is to be solved in such cases to obtain the displacement fluctuations in the interior of the RVE. The proviso (3.1.5) is satisfied, for instance, with *zero* boundary fluctuations ($\mathbf{u}_\mu(\mathbf{y}, t) = \mathbf{0}, \forall \mathbf{y} \in \partial\Omega_\mu$), and, less trivially, with appropriately prescribed *periodic* boundary conditions.

3.1.3 Homogenized stress

By virtue of assumption (3.1.1), we have that

$$\delta\boldsymbol{\epsilon}_\mu = \delta\boldsymbol{\epsilon} + \nabla^s \delta\mathbf{u}_\mu. \quad (3.1.6)$$

Substitution of this equation in the variational statement of the Hill-Mandell principle of macro-homogeneity (Eq.(3.1.4)) leads to

$$\boldsymbol{\sigma} : \delta\boldsymbol{\epsilon} = \frac{1}{V_\mu} \int_{\Omega_\mu} \boldsymbol{\sigma}_\mu : \delta\boldsymbol{\epsilon} d\Omega + \frac{1}{V_\mu} \int_{\Omega_\mu} \boldsymbol{\sigma}_\mu : \nabla^s \delta\mathbf{u}_\mu d\Omega. \quad (3.1.7)$$

The above equation must be obeyed for any kinematically admissible $\delta\mathbf{u}_\mu$. In particular, if $\delta\mathbf{u}_\mu = \mathbf{0}$, we get that

$$\boldsymbol{\sigma} : \delta\boldsymbol{\epsilon} = \frac{1}{V_\mu} \int_{\Omega_\mu} \boldsymbol{\sigma}_\mu : \delta\boldsymbol{\epsilon} d\Omega, \quad \forall \delta\boldsymbol{\epsilon}. \quad (3.1.8)$$

Hence,

$$\boldsymbol{\sigma} = \frac{1}{V_\mu} \int_{\Omega_\mu} \boldsymbol{\sigma}_\mu d\Omega, \quad (3.1.9)$$

i.e., the macroscopic stress tensor $\boldsymbol{\sigma}$ is the volume average of the microscopic stress field $\boldsymbol{\sigma}_\mu = \boldsymbol{\sigma}_\mu(\mathbf{y})$, $\mathbf{y} \in \Omega_\mu$ over the associated RVE².

Another condition that emerges from the variational equation (3.1.7) is that

$$\int_{\Omega_\mu} \boldsymbol{\sigma}_\mu : \nabla^s \delta\mathbf{u}_\mu d\Omega, \quad (3.1.10)$$

for any $\delta\mathbf{u}_\mu \in \mathcal{V}_\mu$. It can be shown (de Souza Neto and Feijóo, 2006) that this condition amounts to requiring that the external surface traction and body force field in the RVE be purely reactive —i.e., a reaction to the kinematical constraints imposed upon the RVE. This is why the RVE equilibrium equation presented in the next section does not contain external boundary traction nor body force terms.

3.1.4 Incremental variational formulation of the equilibrium of the RVE

Consider a time discretization of the interval of interest $[t_0, t_f] = \bigcup_{n=1}^N [t_n, t_{n+1}]$. We shall presume that the current value of the microscopic stress tensor $\boldsymbol{\sigma}_\mu|_{n+1}$ at each $\mathbf{y} \in \Omega_\mu^s$ is entirely determined by the current value of the microscopic strain tensor $\boldsymbol{\epsilon}_{\mu_{n+1}}(\mathbf{y}) = \boldsymbol{\epsilon}_{n+1} + \nabla^s \mathbf{u}_\mu(\mathbf{y})$, on the one hand, and a set of microscopic internal variables $\boldsymbol{\xi}_\mu|_{n+1}$, on the other hand —that encapsulates the history of microscopic deformations. The relationship between these variables is established by (phenomenological) incremental constitutive equations; for multiphase materials, these constitutive equations may vary from point to point within the RVE. The (incremental) RVE equilibrium problem at time t_{n+1} can be stated as follows. Given the *initial*

²Thus, the macroscopic stress tensor $\boldsymbol{\sigma}$ being the volume average of the microscopic stress field $\boldsymbol{\sigma}_\mu = \boldsymbol{\sigma}_\mu(\mathbf{y})$, $\mathbf{y} \in \Omega_\mu$ over the associated RVE is a corollary of the Hill-Mandell principle and the additive decomposition of the microscopic strain, rather than, as often erroneously claimed, an axiomatic presumption of the employed homogenized approach.

data $\{\mathbf{u}_\mu|_n(\mathbf{y}), \boldsymbol{\epsilon}_n, \boldsymbol{\xi}_\mu|_n(\mathbf{y})\}$ and the *prescribed* macroscopic strain tensor $\boldsymbol{\epsilon}_{n+1}$, find a displacement fluctuations field $\mathbf{u}_\mu|_{n+1}(\mathbf{y}) \in \mathcal{V}_\mu$ such that:

$$\hat{G}_{n+1} = \int_{\Omega_\mu^s} \boldsymbol{\sigma}_\mu(\boldsymbol{\epsilon}_{n+1}, \mathbf{u}_\mu|_{n+1}) : \nabla^s \boldsymbol{\eta} d\Omega - \int_{\partial\Omega_\mu^v} \mathbf{t}_{n+1}^v \cdot \boldsymbol{\eta} d\Gamma = 0, \quad (3.1.11)$$

for all³ $\boldsymbol{\eta} \in \mathcal{V}_\mu$. Here, \hat{G}_{n+1} stands for the incremental virtual work functional at t_{n+1} , $\boldsymbol{\eta}$ denotes the test function, and \mathbf{t}_{n+1}^v represents the *internal traction* exerted upon the solid part of Ω_μ across the solid-void interface $\partial\Omega_\mu^v$ —due, for instance, to the presence of pressurized fluid phase within the voids. In this work, attention is restricted to porous materials in the presence of only solid phases (empty pores); the contribution of internal tractions, thus, vanishes, and Eq.(3.1.11) boils down to:

$$\hat{G}_{n+1} = \int_{\Omega_\mu^s} \boldsymbol{\sigma}_\mu(\boldsymbol{\epsilon}_{n+1}, \mathbf{u}_\mu|_{n+1}) : \nabla^s \boldsymbol{\eta} d\Omega = 0, \quad \forall \boldsymbol{\eta} \in \mathcal{V}_\mu. \quad (3.1.12)$$

The superindex “n+1” will be hereafter drop out, and quantities will be assumed to be evaluated at time t_{n+1} ; only when confusion is apt to show up, the pertinent distinction will be introduced.

3.2 Finite element discretization: the full-order model (FOM)

Following common finite element procedure, the solution of the above problem is sought in the finite dimensional space $\mathcal{V}_\mu^h \in \mathcal{V}_\mu$ spanned by the linearly independent functions $\{N_1, N_2 \dots N_n\}$ (n denotes the number of nodes of the discretization). Accordingly, problem 3.1.12 is now posed in \mathcal{V}_μ^h , and $\boldsymbol{\eta}$ and \mathbf{u}_μ^h are expressed as linear combinations of the basis functions:

$$\mathbf{u}_\mu^h = \sum_{I=1}^n \hat{\mathbf{u}}_{\mu,I}^h N_I(\mathbf{y}), \quad (3.2.1)$$

$$\boldsymbol{\eta}^h = \sum_{I=1}^n \hat{\boldsymbol{\eta}}_I^h N_I(\mathbf{y}). \quad (3.2.2)$$

Inserting the above approximations in Eq.(3.1.12), and exploiting the arbitrariness of the coefficients $\hat{\boldsymbol{\eta}}_I^h$, we get:

$$R_{iI} = \int_{\Omega_\mu^s} \frac{\partial N_I}{\partial x_j} (\boldsymbol{\sigma}_\mu)_{ji} d\Omega = 0 \quad (i = 1 \dots d; I = 1 \dots n), \quad (3.2.3)$$

where R_{iI} stands for the i – th component of the residual at the I – th node. Introducing the classical “ \mathbf{B} -matrix” connecting strain tensor with nodal displacements⁴:

$$(\nabla^s \mathbf{u}_\mu^h)_{ij} = \text{sym} \left(\frac{\partial N_I}{\partial x_j} \delta_{ki} \right) (\hat{\mathbf{u}}_{\mu,I}^h)_k = (\mathbf{B})_{ijIk} (\hat{\mathbf{u}}_{\mu,I}^h)_k, \quad (3.2.4)$$

³Notice that we are tacitly assuming that the space of trial and test functions coincide (*zero boundary fluctuations*). We should emphasize that this is done solely for simplicity in the upcoming finite element formulation; the procedure for constructing reduced basis approximations is the same for both zero and periodic boundary conditions—in fact, for any model with boundary conditions consistent with proviso Eq.(3.1.5).

⁴Index notation used follows Belytschko et al. (2001)

eq. (3.2.3) can be rewritten as

$$R_{iI} = \int_{\Omega_\mu^s} (\mathbf{B})_{rjIi} (\boldsymbol{\sigma}_\mu)_{jr} d\Omega = 0 \quad (i = 1 \dots d; I = 1 \dots n), \quad (3.2.5)$$

or in matrix form (Voigt's notation)⁵:

$$\mathbf{R} = \int_{\Omega_\mu^s} \mathbf{B}^T \{ \boldsymbol{\sigma}_\mu \} d\Omega = \mathbf{0}, \quad (3.2.6)$$

where the brackets around $\boldsymbol{\sigma}_\mu$ indicates that the stress tensor is stored in matrix form.

Needless to say, the dependence of $\boldsymbol{\sigma}_\mu$ on the nodal displacements is, in general, non-linear, and the solution of Eq.(3.2.6) requires an iterative solution. The classical Newton-Raphson iterative scheme is constructed by applying to Eq.(3.2.6) the Gateaux derivative in the direction of an incremental fluctuation $\Delta \mathbf{u}_\mu^h$ (denoted by $\mathcal{D}(\bullet)[\Delta \mathbf{u}_\mu^h]$):

$$R_{iI}|^{(k+1)} = 0 = R_{iI}|^{(k)} + \int_{\Omega_\mu^s} (\mathbf{B})_{rjIi} \mathcal{D}(\boldsymbol{\sigma}_\mu)_{jr} [\Delta \mathbf{u}_\mu^h] d\Omega \quad (3.2.7)$$

where the superindex “(k)” is the iteration counter. The rightmost term in the above is expressible as:

$$\begin{aligned} \int_{\Omega_\mu^s} (\mathbf{B})_{rjIi} \mathcal{D}(\boldsymbol{\sigma}_\mu)_{jr} [\Delta \mathbf{u}_\mu] d\Omega &= \int_{\Omega_\mu^s} (\mathbf{B})_{rjIi} (\mathbf{C}_\mu)_{jrst} (\nabla^s \Delta \mathbf{u}_\mu^h)_{st} d\Omega \\ &= \int_{\Omega_\mu^s} (\mathbf{B})_{rjIi} (\mathbf{C}_\mu)_{jrst} (\mathbf{B})_{stMu} \Delta \hat{\mathbf{u}}_{\mu,uM}^h d\Omega \\ &= \int_{\Omega_\mu^s} \overbrace{(\mathbf{K}_\mu)_{iIMu}}^{(\mathbf{K}_\mu)_{iIMu}} (\mathbf{B})_{rjIi} (\mathbf{C}_\mu)_{jrst} (\mathbf{B})_{stMu} d\Omega \Delta \hat{\mathbf{u}}_{\mu,uM}^h \\ &= (\mathbf{K}_\mu)_{iIMu} \Delta \hat{\mathbf{u}}_{\mu,uM}^h, \end{aligned} \quad (3.2.8)$$

where \mathbf{C}_μ is the *algorithmic tangent operator* consistent with the microscopic incremental constitutive law, and \mathbf{K}_μ denotes the tangent stiffness matrix; in Voigt's notation, \mathbf{K}_μ can be written as:

$$\{ \mathbf{K}_\mu \} = \int_{\Omega_\mu^s} \{ \mathbf{B}^T \} \{ \mathbf{C}_\mu \} \{ \mathbf{B} \} d\Omega. \quad (3.2.9)$$

In turn, with Eq.(3.2.9) at our disposal, Eq.(3.2.7) can be also rephrased in Voigt's notation:

$$\mathbf{0} = \{ \mathbf{R} \}^{(k)} + \{ \mathbf{K}_\mu \}^{(k)} \left(\left\{ \hat{\mathbf{u}}_\mu^h \right\}^{(k+1)} - \left\{ \hat{\mathbf{u}}_\mu^h \right\}^{(k)} \right). \quad (3.2.10)$$

⁵Voigt's conversion rules for the pair of indices rj and Ii can be also consulted in Belytschko et al. (2001).

3.2.1 Elastic micro-constitutive law

It proves instructive at this point to particularize the expression derived above to the elastic case, in which $\boldsymbol{\sigma}_\mu = \mathbf{C}_\mu : (\boldsymbol{\epsilon} + \nabla^s \mathbf{u}_\mu)$. Upon trivial manipulation, one arrive at the following equilibrium equation:

$$\int_{\Omega_\mu^s} \overbrace{\{\mathbf{B}^T\} \{\mathbf{C}_\mu\}}^{\mathbf{f}_\epsilon} d\Omega \{\boldsymbol{\epsilon}\} + \{\mathbf{K}_\mu\} \{\hat{\mathbf{u}}_\mu^h\} = \mathbf{0}. \quad (3.2.11)$$

Note that, in the above equation, the first term—the one in which the macrostrain tensor $\boldsymbol{\epsilon}$ appears— plays the role of “external” force; in the case we are considering (homogeneous essential boundary conditions), the absence of this term would imply that the solution to the equilibrium problem is a vanishing fluctuation field. Note also that $\mathbf{K}_\mu \in \mathbb{R}^{N \times N}$, $\mathbf{B} \in \mathbb{R}^{d_s \times N}$, $\hat{\mathbf{u}}_{\mu, \epsilon}^h \in \mathbb{R}^N$ and $\boldsymbol{\epsilon} \in \mathbb{R}^{d_s}$, being $N = n \cdot d$ the number of degrees of freedom (d is the dimension of the problem) and d_s the number of (independent) components of the tensor $\boldsymbol{\epsilon}$.

The solution for Eq.(3.2.11) can be explicitly obtained by, first, partitioning the set of nodes into interior ($\{\hat{\mathbf{u}}_\mu^h\}_{int}$) and boundary ($\{\hat{\mathbf{u}}_\mu^h\}_{bnd}$) nodes:

$$\{\hat{\mathbf{u}}_\mu^h\}^T = \left[\begin{array}{c} \{\hat{\mathbf{u}}_\mu^h\}_{int}^T \\ \{\hat{\mathbf{u}}_\mu^h\}_{bnd}^T \end{array} \right]. \quad (3.2.12)$$

Substituting Eq.(3.2.12) into Eq.(3.2.11), and solving for $\{\hat{\mathbf{u}}_\mu^h\}_{int}$ for the homogeneous boundary conditions case ($\{\hat{\mathbf{u}}_\mu^h\}_{bnd} = \mathbf{0}$) yields:

$$\{\hat{\mathbf{u}}_\mu^h\} = \left[\begin{array}{c} \{\hat{\mathbf{u}}_\mu^h\}_{int} \\ \{\hat{\mathbf{u}}_\mu^h\}_{bnd} \end{array} \right] = \left[\begin{array}{c} \{-\mathbf{K}_\mu\}_{int,int}^{-1} \mathbf{f}_\epsilon \\ \mathbf{0} \end{array} \right] \{\boldsymbol{\epsilon}\} = \mathcal{P}_\epsilon \{\boldsymbol{\epsilon}\}. \quad (3.2.13)$$

The above result concurs with one’s intuitive expectations: in the elastic range, nodal micro-displacements bear a linear relationship with the macro-strain vector $\boldsymbol{\epsilon}$, being \mathcal{P}_ϵ the corresponding projection operator.

3.3 Construction of the reduced-order model (ROM)

3.3.1 Sampling of the parametric space

The basic input variable in the problem under consideration is the macro-strain tensor $\boldsymbol{\epsilon}$; the solution space⁶ is, thus, the set of all nodal displacement fluctuations satisfying Eq.(3.2.6) for different values of $\boldsymbol{\epsilon}$. The first step in constructing a reduced order model is to find an orthonormal set of basis functions for the displacement solution space. This task is carried out by applying the singular value decomposition (see section A) to the so-called *snapshot matrix*.

3.3.1.1 Elastic case

Let us consider first the most simple case: an elastic micro-constitutive law (discussed in section 3.2.1). To obtain orthonormal bases, we evaluate

⁶In fact, in such a parametric space, each vector is a vector-valued function of the time.

Eq.(3.2.13) at N_p different values of the macro-strain tensor, and collect the outcome of such calculations in a single $N \times N_p$ matrix — the snapshot matrix:

$$\mathbf{X}_U^h = [\hat{\mathbf{u}}_\mu^h(\boldsymbol{\epsilon}_1) \quad \hat{\mathbf{u}}_\mu^h(\boldsymbol{\epsilon}_2) \quad \cdots \quad \hat{\mathbf{u}}_\mu^h(\boldsymbol{\epsilon}_{N_p})], \quad (3.3.1)$$

Then, the *singular value decomposition* is applied to \mathbf{X}_U^h :

$$\mathbf{X}_U^h = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^T, \quad (3.3.2)$$

where $\mathbf{U} \in \mathbb{R}^{N \times r_{snap}^u}$ — r_{snap}^u is the rank of \mathbf{X}_U^h — satisfies $\mathbf{U}^T\mathbf{U} = \mathbf{I}$; $\mathbf{V} \in \mathbb{R}^{r_{snap}^u \times N}$ fulfills $\mathbf{V}^T\mathbf{V} = \mathbf{I}$; and $\boldsymbol{\Sigma} = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_{r_{snap}^u})$, with $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_{r_{snap}^u} > 0$. The columns of \mathbf{U} are called *left singular vectors* of \mathbf{X}_U^h , and constitute a set of orthonormal bases for the range of \mathbf{X}_U^h ; the columns of \mathbf{V} , on the other hand, are the *right singular vectors*, and can be regarded as orthonormal bases for the nullspace of \mathbf{X}_U^h ; finally, the scalars σ_i are the *singular values* of \mathbf{X}_U^h . By virtue of Eq.(3.2.13), Eq.(3.3.1) can be written, in the elastic case, as:

$$\mathbf{X}_U^h = \mathcal{P}_\epsilon [\boldsymbol{\epsilon}_1 \quad \boldsymbol{\epsilon}_2 \quad \cdots \quad \boldsymbol{\epsilon}_{N_p}] = \mathcal{P}_\epsilon [\boldsymbol{\epsilon}]_{snap}, \quad (3.3.3)$$

where $[\boldsymbol{\epsilon}]_{snap} \in \mathbb{R}^{d_s \times N_p}$ and $\mathcal{P}_\epsilon \in \mathbb{R}^{N \times d_s}$. It follows — from the property that $\text{rank}(AB) \leq \min(\text{rank}(A), \text{rank}(B))$ — that, in the elastic case, the rank of \mathbf{X}_U^h cannot be greater than d_s . Thus, for purposes of constructing orthogonal bases for the input parametric space, it suffices to calculate nodal micro-displacement fluctuations at d_s linearly independent macro-strain vectors. This means also that the number of non-zero entries on the diagonal of $\boldsymbol{\Sigma}$ will be less or equal than d_s .

3.3.1.2 Inelastic case

As may be surmised, when the constitutive relation between (micro) stress and strains is not linear, the situation is not so clear-cut; as opposed to the elastic case, it is virtually impossible to ascertain *a priori* how many simulations are needed to obtain a set of basis vectors able to cover satisfactorily the input parametric space. Nevertheless, the conclusions reached in the elastic case suggests somehow that the number of necessary basis vectors will depend *only weakly* on the number of degrees of freedom N of the microcell.

Sampling the parametric space efficiently appears, therefore, as a significant challenge⁷. One route to accomplish this task may be to consider as basic input variable strain histories rather than total strains; i.e., simulations are conducted with various strain histories, combinations of compression, extension, shearing, relaxation, etc. The results stored in the snapshot matrix

⁷“... if the dimension of the parameter space is large, uniform sampling will quickly become too computationally expensive due to the combinatorial explosion of samples needed to cover the parameter space” (Bui-Thanh et al., 2007; Bui-Thanh, 2007). For further information on sparse, efficient, sampling strategies, the reader is referred to Bui-Thanh et al. (2008); they address the challenge of sampling a high-dimensional parameter space by exploiting the *greedy sampling method*. The key premise of greedy sampling is to *adaptively choose samples* by finding the location at which the estimate of the error in the reduced model is maximum, over a predetermined discrete set of parameters. The sampling issue is also addressed in Carlberg and Farhat (2008)

are not the “final” — at the end of the strain trajectory — nodal displacements, but rather the evolution of such variables (sampled at discrete time intervals) for the various strain histories.

3.3.2 Galerkin projection onto the reduced subspace

Suppose now that, by means of the SVD of the displacement snapshot matrix \mathbf{X}_U^h , we know that the leading N_u singular values in Σ are significantly greater than the others; or, put it alternatively, that the magnitude of the trailing $r_{snp}^u - N_u$ singular values are comparatively negligible. The matrix formed by the first N_u columns of \mathbf{U} will be henceforth denoted as $\Phi \in \mathbb{R}^{N \times N_u}$.

We now seek to pose the boundary-valued problem represented by Eq.(3.1.12) in the reduced⁸ configuration subspace $\mathcal{V}_\mu^I \subseteq \mathcal{V}_\mu^h$ spanned⁹ by the basis functions defined as¹⁰:

$$\Phi_{\hat{J}}(\mathbf{y}) = \sum_{I=1}^n \Phi_{\hat{J}}^I N_I(\mathbf{y}), \quad \hat{J} = 1, 2, \dots, N_u, \quad (3.3.4)$$

where the *nodal parameters* $\Phi_{\hat{J}}^I$, arranged in column form, are precisely the \hat{J} -th column¹¹ of Φ . To this end, we define a new set of displacement fluctuations and test functions: $\mathbf{u}_\mu^I \in \mathcal{V}_\mu^I$ and $\eta^I \in \mathcal{V}_\mu^I$, defined as

$$\mathbf{u}_\mu^I(\mathbf{y}) = \sum_{\hat{I}=1}^{N_u} \hat{\mathbf{u}}_\mu^I|_{\hat{I}} \Phi_{\hat{I}}(\mathbf{y}) = \sum_{\hat{I}=1}^{N_u} \sum_{J=1}^n \hat{\mathbf{u}}_\mu^I|_{\hat{I}} \Phi_{\hat{I}}^J N_J, \quad (3.3.5)$$

$$\eta^I = \sum_{\hat{I}=1}^{N_u} \hat{\eta}_{\hat{I}}^I \Phi_{\hat{I}} = \sum_{\hat{I}=1}^{N_u} \sum_{J=1}^n \hat{\eta}_{\hat{I}}^I \Phi_{\hat{I}}^J N_J, \quad (3.3.6)$$

Inserting the above approximations in Eq.(3.1.12), exploiting the arbitrariness of $\hat{\eta}_{\hat{I}}^I$, and after some algebra, the following expression (in Voigt's notation) for the residual is obtained:

$$\mathbf{R}^I = \Phi^T \int_{\Omega_\mu^s} \mathbf{B}^T \{ \boldsymbol{\sigma}_\mu(\boldsymbol{\epsilon}, \hat{\mathbf{u}}_\mu^I) \} d\Omega = \mathbf{0}. \quad (3.3.7)$$

Note that $\Phi \in \mathbb{R}^{N \times N_u}$, and, thus, $\mathbf{R}^I \in \mathbb{R}^{N_u}$. The counterpart of Eq.(3.2.10) is constructed in a similar fashion:

$$\mathbf{0} = \{ \mathbf{R}^I \}^{(k)} + \{ \mathbf{K}_\mu^I \}^{(k)} \left(\{ \hat{\mathbf{u}}_\mu^I \}^{(k+1)} - \{ \hat{\mathbf{u}}_\mu^I \}^{(k)} \right). \quad (3.3.8)$$

where $\{ \mathbf{K}_\mu^I \} \in \mathbb{R}^{N_u \times N_u}$ stands for the *reduced stiffness matrix*, defined as:

$$\{ \mathbf{K}_\mu^I \} = \Phi^T \{ \mathbf{K}_\mu \} \Phi = \Phi^T \left(\int_{\Omega_\mu^s} \{ \mathbf{B}^T \} \{ \mathbf{C}_\mu \} \{ \mathbf{B} \} d\Omega \right) \Phi. \quad (3.3.9)$$

⁸A semantic remark is in order here. Whereas the term “discretization” is used for referring to the transition from the infinite dimensional space \mathcal{V}_μ to the finite (element) space \mathcal{V}_μ^h , the term “reduction” connotes a transition from the finite element space \mathcal{V}_μ^h to the sensibly smaller subspace \mathcal{V}_μ^I .

⁹Note that in the elastic case, the two spaces may coincide

¹⁰Strictly, this equality only holds for homogeneous boundary conditions

¹¹Note that the basis functions $\Phi_{\hat{J}}(\mathbf{y})$ fall within the category of Ritz functions — they are globally supported functions, as distinct from standard FE basis functions, which have compact-support (Krysl et al., 2001).

Nevertheless, computing $\{\mathbf{K}_\mu^I\}$ by reducing — via the operator $\Phi^T(\bullet)\Phi$ — the full assembled tangent stiffness matrix $\{\mathbf{K}_\mu\}$ is not computationally efficient; it proves more advantageous to assemble the reduced element matrices, denoted by ${}^e\Phi$:

$$\begin{aligned} \{\mathbf{K}_\mu^I\} &= \Phi^T \left(\int_{\Omega_\mu^s} \{\mathbf{B}^T\} \{\mathbf{C}_\mu\} \{\mathbf{B}\} d\Omega \right) \Phi \\ &= \mathbf{A}_{e=1}^{n_e} {}^e\Phi^T \left(\int_{{}^e\Omega_\mu^s} \{{}^e\mathbf{B}^T\} \{\mathbf{C}_\mu\} \{{}^e\mathbf{B}\} d\Omega \right) {}^e\Phi \quad (3.3.10) \\ &= \mathbf{A}_{e=1}^{n_e} {}^e\Phi^T \{{}^e\mathbf{K}_\mu\} {}^e\Phi. \end{aligned}$$

At first encounter, the fact that the dimensions of the tangent stiffness operator on the reduced basis are $N_u \times N_u$, i.e., they only depend on the cardinality of the reduced basis, may lead to think that the complexity of the problem is only N_u -dependent. Unfortunately, this is not so: the cost to evaluate both the internal forces, which, in turn, involve the update of the stress tensor $\boldsymbol{\sigma}_\mu$ (sometimes the most computationally costly operation), and the element-wise calculation of the stiffness matrix of the FE model remains dependent upon the size (N) of the original model (they scale with N). In fact, the only advantage that accrues in reducing the dimensionality of the model is the solution of the linear system of equations at each Newton-Raphson iteration; the other operations are even more computationally laborious.

3.3.3 Analysis of approximation errors

3.3.3.1 Sampling and truncation errors

Discrepancies between full-order and reduced-order responses arise from two distinct sources, namely: truncation and sampling. The *truncation error* is the error introduced as a result of using a small number of modes; accordingly, this error vanishes as the number of selected modes approaches the rank of the snapshot matrix. On the other hand, the *sampling error* is the manifestation of the fact that the snapshots are but realizations of the state field — the micro-displacement fluctuations — for a (presumably) representative *sample* of points in the parametric space — the set of all physically conceivable strain histories. If the strain histories chosen for generating the snapshots fail notably to represent such a parametric space, the magnitude of the sampling error will be significant. Likewise, the sampling error should vanish identically when the prescribed strain history pertains to the set of strain trajectories employed to calculate the snapshots.

The sampling error can be evaluated for an arbitrary strain trajectory $\boldsymbol{\epsilon}_t$ as the difference between the FOM and ROM solutions in the limit of no truncation ($N_u = r_{snp}^u$). In terms of displacements, thus, this error can be defined formally as follows:

$$\mathbf{e}_{u_\mu}^{I,samp} := \mathbf{u}_\mu^h - \mathbf{u}_\mu^I(r_{snp}^u). \quad (3.3.11)$$

On the other hand, the error attributed to truncation of the basis to the N_u -th mode can be calculated by subtracting the above defined sampling error to the total error:

$$\mathbf{e}_{u_\mu}^{I, \text{trun}}(N_u) := (\mathbf{u}_\mu^h - \mathbf{u}_\mu^I(N_u)) - \mathbf{e}_{u_\mu}^{I, \text{samp}} = \mathbf{u}_\mu^I(r_{\text{snap}}^u) - \mathbf{u}_\mu^I(N_u). \quad (3.3.12)$$

The truncation error, thus, measures the discrepancies between two ROM solutions, and hence does not depend actually upon the FOM solution. Furthermore, it follows easy from its definition that¹² $\mathbf{e}_{u_\mu}^{I, \text{trun}} \in \text{span}\{\Phi_i\}_{i=N_u+1}^{r_{\text{snap}}^u}$.

Similar definitions can be established for the micro-stress field (the actual output of interest in the multiscale problem). In analogy with Eq.(3.3.11), the sampling error in the approximation of the micro-stress field is defined as

$$\mathbf{e}_\Sigma^{I, \text{samp}} := \Sigma^h - \Sigma^I(r_{\text{snap}}^u), \quad (3.3.13)$$

where $\Sigma^h \in \mathbb{R}^{N_g \times d_s}$ denotes the FOM stresses and $\Sigma^I \in \mathbb{R}^{N_g \times d_s}$ stands for the ROM solution in the limit of no truncation ($N_u = r_{\text{snap}}^u$). Likewise, the stress error incurred by truncation of the basis after the $N_u + 1$ vector is given by

$$\mathbf{e}_\Sigma^{I, \text{trun}}(N_u) := \mathbf{e}_\Sigma^I - \mathbf{e}_\Sigma^{I, \text{samp}} = \Sigma^I(r_{\text{snap}}^u) - \Sigma^I(N_u). \quad (3.3.14)$$

3.3.3.2 Consistency

According to Carlberg et al. (2011), the construction of a reduced-order model should be accomplished so that the requirements of *consistency* and *optimality* are fulfilled. An approximation is said to be *consistent* if, when implemented without data compression (no truncation of the basis), it introduces no additional error in the solution of the same problem for which data was acquired. By definition, the truncation error $\mathbf{e}_{u_\mu}^{I, \text{trun}}$ (see Eq.(3.3.12)) vanishes in the limiting case of untruncated basis; thus, consistency is met if the sampling error defined in Eq.(3.3.11) becomes zero when the strain history used for comparison purposes is the same as the one employed for generating the snapshots ($\boldsymbol{\epsilon}_t = \boldsymbol{\epsilon}_t^{\text{snap}}$). It is not difficult to show that, under the hypothesis of uniqueness of the finite element solution, the projection-based reduced-order model is consistent. Indeed, when $\boldsymbol{\epsilon}_t = \boldsymbol{\epsilon}_t^{\text{snap}}$, both FOM and ROM pertains to $\text{span}\{\Phi_i\}_{i=1}^{r_{\text{snap}}^u}$. Lack of consistency would imply that $\mathbf{u}_\mu^h \neq \mathbf{u}_\mu^I$, which is not possible since it would violate the uniqueness¹³ assumption.

¹²It is worthy to note that $\mathbf{e}_{u_\mu}^{I, \text{samp}}$ and $\mathbf{e}_{u_\mu}^{I, \text{trun}}$ are not necessarily orthogonal; $\mathbf{e}_{u_\mu}^{I, \text{samp}}$ can be further resolved into two mutually orthogonal components: one that lies in $\text{span}\{\Phi_i\}_{i=1}^{r_{\text{snap}}^u}$, and another that resides in the orthogonal complement of $\text{span}\{\Phi_i\}_{i=1}^{r_{\text{snap}}^u}$. It would be interesting in future research to see how the sampling error is apportioned between these two complementary spaces. Intuitively, one would expect that the component in the orthogonal complement of the space spanned by the snapshots dominates. This component is defined as $\mathbf{e}_{u_\mu}^{I, \text{samp}^\perp} = (\mathbf{I} - \Phi\Phi^T)\mathbf{u}_\mu^h$. It transpires somehow that the other component — the one that $\text{span}\{\Phi_i\}_{i=1}^{r_{\text{snap}}^u}$ — is due to the macro-strain path-dependent nature of the problem.

¹³Consistency may, thus, not be observed in cases in which the stress-deformation relation is not uniquely invertible (strain softening).

3.3.3.3 Optimality

In regard to the optimality requirement, according to Carlberg et al. (2011), “an approximation is said to be optimal if it leads to approximated quantities that minimize some error measure”. Carlberg et al. (2011) argue that optimality should be reflected in a monotonic decay of some norm the truncation error as the subspace spanned by the basis expands. By construction, the POD basis leads naturally to optimal approximations¹⁴ in the state variable¹⁵ (micro-displacement field). Unfortunately, optimality cannot be guaranteed for other relevant variables in the problem, in particular, the micro-stress field (the actual output of interest). To optimize the selected displacement modes with respect to a particular output functional, recourse is to be made to more sophisticated strategies such as the *goal-oriented approach* proposed by¹⁶ Bui-Thanh et al. (2007).

¹⁴The central idea of POD is, as pointed out by Rowley et al. (2004), “to compute a nested family of subspaces, of increasing dimension, that optimally spans the [snapshot] data, in the sense that the error in the projection onto each subspace is minimized”

¹⁵An additional proviso is that the pertinent projection technique should furnish optimal approximations. The Galerkin projection (see Carlberg et al. (2011)) is not optimal when the jacobians are not semi-definite positive; in such a situation, recourse is to be made to Petrov-Galerkin projections.

¹⁶A key aspect in the construction of a goal-oriented approach is the appropriate choice of the *inner product* used for defining the minimization problem that leads to the POD basis. Although in this work it has been tacitly assumed that this inner product is the standard euclidean inner product, the reader is to be aware that other choices are possible. Rowley et al. (2004) introduce an interesting distinction between energy-based and non-energy-based inner products. According to Rowley et al. (2004), an energy-based inner product is one in which the energy defined by the induced norm is a meaningful physical quantity. Clearly, the euclidean inner product for the problem at hand (displacement field as the state variable) does not fall into this category. Should the state variable be the velocity, and the inner product the L2-norm, the energy induced would be connected with the kinetic energy.

Chapter 4

Hyperreduced-order modeling (HRoM) of the RVE

4.1 Matrix formulation

To address the issue of the efficient evaluation of the nonlinear terms in the governing equations, it proves convenient first to rephrase equation (3.3.7) for the residual. By defining the “reduced B-matrix” $\mathbf{B}^I \in \mathbb{R}^{d_s \times N_u}$ as:

$$\mathbf{B}^I = \mathbf{B}\Phi, \quad (4.1.1)$$

Eq.(3.3.7) can be rewritten as:

$$\mathbf{R}^I = \int_{\Omega_\mu^s} \mathbf{B}^{IT} \{\boldsymbol{\sigma}_\mu\} d\Omega = \mathbf{0}. \quad (4.1.2)$$

Now the integral above is approximated by gauss quadrature:

$$\mathbf{R}^I = \sum_{g=1}^{N_g} w_g \mathbf{B}_g^{IT} \{\boldsymbol{\sigma}_\mu\}_g = \mathbf{0}, \quad (4.1.3)$$

where¹ N_g is the total number of gauss points in the mesh, w_g denotes the weight² at the g -th quadrature point, and \mathbf{B}_g^I and $\{\boldsymbol{\sigma}_\mu\}_g$ are the reduced B-matrix and the stress vector evaluated at such gauss point. For later purposes, it is conceptually — and symbolically — advantageous to define a column vector $\boldsymbol{\Sigma}^I \in \mathbb{R}^{N_g \cdot d_s}$ containing the stress vectors at each gauss point:

$$\boldsymbol{\Sigma}^I = \left[\{\boldsymbol{\sigma}_\mu\}_1^T \quad \{\boldsymbol{\sigma}_\mu\}_2^T \quad \cdots \quad \{\boldsymbol{\sigma}_\mu\}_g^T \quad \cdots \quad \{\boldsymbol{\sigma}_\mu\}_{N_g}^T \right]^T, \quad (4.1.4)$$

With the above definition at hand, Eq.(4.1.3) can be legitimately rewritten in the following matrix format:

$$\mathbf{R}^I = \mathbb{B}^{IT} \boldsymbol{\Sigma}^I = \mathbf{0}, \quad (4.1.5)$$

¹Note that in the reduced order model, the the assembly operator degenerates into a simple summation.

²This weight includes both the quadrature weight and the corresponding jacobian determinant.

where $\mathbb{B}^{IT} \in \mathbb{R}^{N_u \times N_g d_s}$ —the generalized strain-displacement B -matrix— is given by:

$$\mathbb{B}^{IT} = \begin{bmatrix} w_1 \mathbf{B}_1^{IT} & w_2 \mathbf{B}_2^{IT} & \cdots & w_g \mathbf{B}_g^{IT} & \cdots & w_{N_g} \mathbf{B}_{N_g}^{IT} \end{bmatrix}. \quad (4.1.6)$$

Note that \mathbb{B}^{IT} can be computed *a priori* — i.e., in the offline stage; the bottleneck that precludes the efficient on-line computations is the calculation of the gauss point stress vector $\boldsymbol{\Sigma}^I$. In the ensuing discussion, we shall attempt to elucidate how to reduce the computational cost associated with the calculation of $\boldsymbol{\Sigma}^I$.

4.2 Approximation of the gauss point stress vector

The strategy followed here to reduce the online computational cost of calculating $\boldsymbol{\Sigma}^I$ at each step is based on the replacement of the non-linear function $\boldsymbol{\Sigma}^I = \boldsymbol{\Sigma}^I(\boldsymbol{\epsilon} + \nabla^s \mathbf{u}_\mu^I)$ by a *coefficient-function* approximation — in the terminology of Nguyen and Peraire (2008) — consisting in a linear combination of *pre-computed* basis and coefficients depending on $\boldsymbol{\epsilon}$. The basis will be constructed by the *method of snapshots*: we collect, in a snapshot matrix \mathbf{X}_σ^h , N_s observations of the vector $\boldsymbol{\Sigma}^I$ for several, representative³ $\boldsymbol{\epsilon}$:

$$\mathbf{X}_\sigma^h = [\boldsymbol{\Sigma}^h(\boldsymbol{\epsilon}_1) \quad \boldsymbol{\Sigma}^h(\boldsymbol{\epsilon}_2) \quad \cdots \quad \boldsymbol{\Sigma}^h(\boldsymbol{\epsilon}_{N_s})], \quad (4.2.1)$$

The basis vectors $\{\boldsymbol{\Psi}_i\}_{i=1\dots}$ are obtained as the first N_σ — the leading (dominant) modes — left singular vectors arising⁴ from the SVD of \mathbf{X}_σ^h .

Next step is to approximate $\boldsymbol{\Sigma}^I$ in Eq.(4.1.5) by a linear combination of the above calculated basis:

$$\boldsymbol{\Sigma}^I(\boldsymbol{\epsilon}, \mathbf{u}_\mu) \approx \boldsymbol{\Sigma}^{II}(\boldsymbol{\epsilon}, \mathbf{u}_\mu) = \sum_{i=1}^{N_\sigma} \boldsymbol{\Psi}_i c_i(\boldsymbol{\epsilon}, \mathbf{u}_\mu), \quad (4.2.2)$$

or in matrix form:

$$\boldsymbol{\Sigma}^{II}(\boldsymbol{\epsilon}, \mathbf{u}_\mu) = \boldsymbol{\Psi} \mathbf{c}(\boldsymbol{\epsilon}, \mathbf{u}_\mu), \quad (4.2.3)$$

where $\boldsymbol{\Psi} \in \mathbb{R}^{N_g d_s \times N_\sigma}$ is the basis matrix and $\mathbf{c}(\in \mathbb{R}^{N_\sigma})$ a certain vector of coefficients. One approach to calculate such a coefficient vector is through the *gappy data reconstruction method* (Everson and Sirovich, 1995); according to such a method, \mathbf{c} is given by

$$\mathbf{c} = \left(\hat{\boldsymbol{\Psi}}^T \hat{\boldsymbol{\Psi}} \right)^{-1} \hat{\boldsymbol{\Psi}}^T \hat{\boldsymbol{\Sigma}}^{II}. \quad (4.2.4)$$

In the above, $\hat{\boldsymbol{\Psi}} \in \mathbb{R}^{\hat{N}_g d_s \times N_\sigma}$ is a sub-block matrix of $\boldsymbol{\Psi}$ constructed by collecting the rows corresponding to $\hat{N}_g \geq N_\sigma$ strategically⁵ selected gauss

³The choice of appropriate parameters $\boldsymbol{\epsilon}$ should be subjected to in-depth study; the quality of the approximation depends crucially on it.

⁴A detailed discussion on the the POD basis construction can be found in appendix A

⁵The choice must ensure the existence of the inverse of $\hat{\mathbf{M}} = \hat{\boldsymbol{\Psi}}^T \hat{\boldsymbol{\Psi}}$ (admissible sample points); see appendix C for a discussion of the various selection (or sampling) algorithms ensuring this desirable feature. By elementary algebra, we know that a necessary condition for this to happen is that the number of selected rows n_{row} must be equal or greater than the number columns of the matrix $\boldsymbol{\Psi}$ ($n_{row} \geq N_\sigma$). However, since stress components are

points — the *gappy* data; and $\hat{\Sigma}^{II} \in \mathbb{R}^{\hat{N}_g d_s}$ denotes the vector that contains the stresses at such gauss points.

The approximated problem can be thus posed as follows: for any ϵ , find $\hat{\mathbf{u}}_\mu^{II} \in \mathbb{R}^{N_u}$ such that:

$$\mathbf{R}^{II} = \mathbb{B}^{IT} \Psi \mathbf{c}(\epsilon, \hat{\mathbf{u}}_\mu^{II}) = 0. \quad (4.2.5)$$

4.3 Ill-posedness of the reduced order problem

Unfortunately, the problem defined by Eq.(4.2.5) is ill-posed. Indeed, if matrix \mathbb{B}^{IT} is regarded as a linear application from $\mathbb{R}^{N_g d_s}$ to \mathbb{R}^{N_u} , then, by virtue of Eq.(4.1.5), any self-equilibrated stress vector Σ^I pertains to the nullspace of \mathbb{B}^{IT} , denoted hereafter by $\text{null } \mathbb{B}^{IT}$. Since the columns of the snapshot matrix are elements of $\text{null } \mathbb{B}^{IT}$, it follows that $\Psi_i \in \text{null } \mathbb{B}^{IT}$; hence:

$$\begin{aligned} \mathbf{R}^{II} &= \mathbb{B}^{IT} \Psi \mathbf{c}(\epsilon, \hat{\mathbf{u}}_\mu^{II}) = \\ &= \sum_i^{N_\sigma} \overbrace{(\mathbb{B}^{IT} \Psi_i)}^{=0} c_i = 0, \quad \text{for any } c_i \in \mathbb{R}; \end{aligned} \quad (4.3.1)$$

i.e., the equation is fulfilled regardless of the value of the displacement fluctuations $\hat{\mathbf{u}}_\mu^{II}$.

Remark 4.3.1 *The fact that any self-equilibrated stress vector Σ^I pertains to $\text{null } \mathbb{B}^{IT}$ has far-reaching implications. Indeed, suppose that the singular value decomposition is carried out now on matrix \mathbb{B}^{IT} :*

$$\mathbb{B}^{IT} = \mathbf{U}_{Bt} \mathbf{S}_{Bt} \mathbf{V}_{Bt}^T. \quad (4.3.2)$$

The columns of \mathbf{V}_{Bt} whose same-numbered elements $(\mathbf{S}_{Bt})_j$ are zero are an orthonormal basis for the nullspace (see appendix A.2). Since

$$\dim(\mathbb{R}^{N_g d_s}) = N_g d_s = \dim(\text{null } \mathbb{B}^{IT}) + \dim(\text{range}(\mathbb{B}^{IT})) \quad (4.3.3)$$

(see (Šolín, 2006, p. 332)), the cardinality of such an orthogonal basis is equal⁶ to $N_g d_s - \text{rank}(\mathbb{B}^{IT})$.

We can go further in this digression connected with the nullspace of \mathbb{B}^{IT} . The set of all solutions of the problem defined by Eq.(4.2.5) is:

$$\text{null } \mathbb{B}^{IT} \cap \text{span } \{\Psi_i\}_i^S. \quad (4.3.4)$$

Obviously, since $\text{span } \{\Psi_i\}_i^S$ is a subspace of $\text{null } \mathbb{B}^{IT}$, then any $\Sigma^I \in \text{span } \{\Psi_i\}_i^S$ is an self-equilibrated stress.

calculated *in tandem* from the corresponding constitutive equations —it makes no sense to calculate stress components individually—, for each row determined by the sampling algorithm, the remaining $d_s - 1$ rows associated to the spatial point will be also included in the sub-block matrix $\hat{\Psi}$; this will surely improve the accuracy of the approximation at negligible extra-cost. This is why we have replaced the necessary condition $n_{row} \geq N_\sigma$ by the less restrictive requirement $\hat{N}_g \geq N_\sigma$ (in the sense that $n_{row} = d_s \hat{N}_g$).

⁶The rank of \mathbb{B}^{IT} must be necessary equal to the number of displacement modes N_u ; otherwise the stiffness matrix of the reduced order problem would have been singular, which is not possible since we are presupposing the reduced-order equations have already been solved without convergence flaws.

4.4 Proposed solutions

4.4.1 First “tentative” reformulation of the reduced-order model

It transpires from the foregoing discussion that the root cause of the ill-posedness is that the subspace spanned by the POD basis vectors Ψ is a subspace of null \mathbb{B}^{IT} , and hence $\mathbb{B}^{IT}\Psi_i = 0$, $i = 1 \dots N_\sigma$. This somehow suggests a tentative solution to this problem: *the non-linear term subjected to the coefficient-function approximation must not reside in the nullspace of the operator \mathbb{B}^{IT}* . How to ensure this condition? Let us first try to add some vector $\mathbf{f} \in \mathbb{R}^{N_u}$ to the left and right sides of (4.1.5):

$$\mathbb{B}^{IT}\Sigma^I + \mathbf{f} = \mathbf{f}. \quad (4.4.1)$$

Defining \mathbf{f} as:

$$\mathbf{f} = \mathbb{B}^{IT}\Sigma_0^I = \mathbb{B}^{IT}\Sigma_0^I(\epsilon), \quad (4.4.2)$$

Eq.(4.1.5) can be rewritten as:

$$\mathbb{B}^{IT}\overbrace{(\Sigma^I + \Sigma_0^I)}^{\Sigma_\eta^I} = \mathbb{B}^{IT}\Sigma_\eta^I = \mathbf{f}(\epsilon). \quad (4.4.3)$$

The approximation is now carried out over the “perturbed” gauss point stress vector Σ_η^I . Denoting by Ψ_η the corresponding basis matrix, the statement of the problem can be redefined as follows: for any ϵ , find $\hat{\mathbf{u}}_\mu^{II} \in \mathbb{R}^{N_u}$ such that:

$$\mathbb{B}^{IT}\Psi_\eta \mathbf{c}(\epsilon, \hat{\mathbf{u}}_\mu^{II}) = \mathbf{f}(\epsilon). \quad (4.4.4)$$

Note that, unlike Eq.(4.2.5), the solution to this equation is — apparently — not trivial, since $\mathbb{B}^{IT}\Psi_{\eta_i} \neq 0$ provided that $\Psi_{\eta_i} \notin \text{null } \mathbb{B}^{IT}$

4.4.1.1 Choice of the perturbing function

It is easy to see that the only condition that the perturbing stress Σ_0^I has to meet is that $\Sigma_0^I \notin \text{span}\{\Psi_i\}_i^S$, or equivalently, $\Sigma_0^I \notin \text{range}(\mathbf{X}_\sigma^I)$ — otherwise it may be expressed as a linear combination of the basis Ψ_i . Any vector pertaining to the *orthogonal complement* of $\text{range}(\mathbf{X}_\sigma^I)$ naturally fulfills this condition⁷; thus, we have to find a vector Σ_0^I such that:

$$\mathbf{X}_\sigma^{IT}\Sigma_0^I = \mathbf{0}, \quad (4.4.5)$$

i.e., Σ_0^I must be in the nullspace of \mathbf{X}_σ^{IT} . By virtue⁸ of what was stated in remark 4.3.1, the left singular vectors arising from the SVD of \mathbf{X}_σ^I whose same-numbered singular values are identically zero constitute an *orthonormal basis for the nullspace of \mathbf{X}_σ^{IT}* ; i.e.:

$$\mathbf{X}_\sigma^{IT} = [\mathbf{V}'_\Sigma \quad \mathbf{V}''_\Sigma] \begin{bmatrix} \Sigma_\Sigma & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} [\mathbf{U}' \quad \Psi^\perp]^T \quad (4.4.6)$$

⁷See (Reddy, 1998, p. 124)

⁸Take into account that the left singular vectors of \mathbf{X}_σ^I are the right singular values of \mathbf{X}_σ^{IT}

where $\Psi^\perp \in \mathbb{R}^{N_g d_s \times N_g d_s - r_{snp}^\sigma}$ (r_{snp}^σ denotes the rank of \mathbf{X}_σ^I). Accordingly, we have a total number of $N_g d_s - r_{snp}^\sigma$ candidates.

However, even if we arrange things to ensure that $\Sigma_0^I \in \text{null } \mathbf{X}_\sigma^{IT}$, the problem remains ill-posed: it is not hard to show that any basis vector $(\Psi_\eta)_i$ can be written as $(\Psi_\eta)_i = (\Psi)_j d_{ji} + \beta_i \Sigma_0^I$. Inserting this into Eq.(4.4.4), we arrive at:

$$\mathbb{B}^{IT} \overbrace{((\Psi)_j d_{ji} + \beta_i \Sigma_0^I)}^{=0} (\mathbf{c})_i = \mathbb{B}^{IT} \Sigma_0^I. \quad (4.4.7)$$

Therefore:

$$\Sigma_0^I \beta_i (\mathbf{c})_i = \mathbb{B}^{IT} \Sigma_0^I. \quad (4.4.8)$$

The above poses a trivial solution independent of the empirical information contained in the snapshot matrix; hence the unfeasibility of this approach.

4.4.2 Second tentative reformulation

In view of the apparently, ill-founded nature of the previous tentative solution, several proposals were generated during intensive brainstorming sessions. One solution that showed to advantage in terms of robustness and accuracy is described in the following.

4.4.2.1 Residual-based perturbation

Suppose that, in solving the first-reduction order problem outlined in section 3.3; i.e., for⁹ each ${}^j \epsilon_{n+1}$ ($j = 1, 2 \dots N_h$; $n = 0, 1, 2 \dots N_t - 1$), find $\hat{\mathbf{u}}_\mu^I|_{n+1} \in \mathbb{R}^{N_u}$ such that:

$$\mathbf{R}_{n+1}^I = \mathbb{B}^{IT} \Sigma^I ({}^j \epsilon_{n+1}, \hat{\mathbf{u}}_\mu^I|_{n+1}) = 0, \quad (4.4.9)$$

we arrange things such that the *residual* at the very first iteration of the Newton-Rapshon algorithm — the algorithm leading to the pertinent equilibrating displacement field at each time step — is conveniently stored in memory. The solution proposed to alleviate the abovementioned ill-posedness is to construct the perturbed basis Ψ_η from the stresses Σ_η^I emerging from the following alternative problem: for each ${}^j \epsilon_{n+1}$, find $\hat{\mathbf{u}}_\mu^I|_{n+1} \in \mathbb{R}^{N_u}$ such that:

$$\mathbf{R}_{\eta, n+1}^I = \mathbb{B}^{IT} \Sigma_\eta^I ({}^j \epsilon_{n+1}, \hat{\mathbf{u}}_\mu^I|_{n+1}) = \eta \mathbf{R}_{n+1, j}^{I(0)}, \quad (4.4.10)$$

where

$$\mathbf{R}_{n+1, j}^{I(0)} = \mathbf{R}^I ({}^j \epsilon_{n+1}, \hat{\mathbf{u}}_\mu^I|_{n+1}^{(0)}) = \mathbb{B}^{IT} \overbrace{\Sigma_{0, n+1}^I}^{(0)} ({}^j \epsilon_{n+1}, \hat{\mathbf{u}}_\mu^I|_{n+1}^{(0)}). \quad (4.4.11)$$

and η denotes a certain “perturbing” factor.

Thus, Eq.(4.4.9) is equivalent to solve:

⁹It proves conceptually convenient to refine our notational scheme at this point. ${}^j \epsilon_{n+1}$ denotes the macro-strain tensor at time t_{n+1} ($n = 1, 2 \dots N_t$) for the j -th strain history (it is assumed, for simplicity, that each strain history is divided into the same number N_t of time-intervals); thus, the micro-cell problem is studied under a total number of $N_t \cdot N_h$ macro-strain instances (N_h stands for the number of strain histories).

$$\mathbb{B}^{IT} \left(\boldsymbol{\Sigma}_{\boldsymbol{\eta}}^I \Big|_{n+1} - \eta \boldsymbol{\Sigma}_{0,n+1}^I \right) = 0. \quad (4.4.12)$$

The above explained proposal proved quite robust and accurate for certain values¹⁰ of the perturbing factor η ; however, the reason behind why this approach provides such desirable properties is not apparent at first counter. Let us try to throw some light on the possible causes for the “goodness” of the residual-based perturbation.

First, note that by virtue of Eq.(4.4.12):

$$\boldsymbol{\Sigma}_{\boldsymbol{\eta}}^I - \eta \boldsymbol{\Sigma}_0^I \in \text{null } \mathbb{B}^{IT}. \quad (4.4.13)$$

thus $\boldsymbol{\Sigma}^I = \boldsymbol{\Sigma}_{\boldsymbol{\eta}}^I - \eta \boldsymbol{\Sigma}_0^I$ is the self-equilibrated stress vector we wish to efficiently compute. Furthermore, since $\boldsymbol{\Sigma}_0^I$ is the stress vector corresponding to the first iteration, then $\mathbb{B}^{IT} \boldsymbol{\Sigma}_0^I \neq 0$; hence

$$\boldsymbol{\Sigma}_0^I \notin \text{null } \mathbb{B}^{IT}. \quad (4.4.14)$$

It follows, thus, that the columns of the snapshot matrix from which the “perturbed” basis $\boldsymbol{\Psi}_{\boldsymbol{\eta}}$ are computed — via the SVD — are constructed as a linear combination of vectors that do pertain to the nullspace of \mathbb{B}^{IT} — the difference $\boldsymbol{\Sigma}_{\boldsymbol{\eta}}^I - \eta \boldsymbol{\Sigma}_0^I$, — and vectors that are not in the nullspace of \mathbb{B}^{IT} — the perturbing vector $\boldsymbol{\Sigma}_0^I$, —, and hence, can be resolved into a component that resides in $\text{null } \mathbb{B}^{IT}$ and another contained in the orthogonal complement of $\text{null } \mathbb{B}^{IT}$, which is actually the range¹¹ of \mathbb{B}^I . It transpires, thus, that the spirit of this approach is somehow the same as the one proposed in section 4.4.1.1: to perturb with elements that belongs to the subspace orthogonal to $\text{null } \mathbb{B}^{IT}$. The key difference is that, in the proposal put forward in section 4.4.1.1, all the columns of snapshot matrix were affected by an *unique* orthogonal vector; by contrast, in the residual-based approach each column of the snapshot matrix is affected by a different perturbing vector $\boldsymbol{\Sigma}_0^I = \boldsymbol{\Sigma}^I(j\boldsymbol{\epsilon}, \hat{\mathbf{u}}_{\mu}^I|^{(0)})$. As a consequence, the rank of $\mathbb{B}^{IT} \boldsymbol{\Psi}_{\boldsymbol{\eta}}$ might result greater than one — even it may be (hopefully) equal¹² to N_u , a necessary condition for the Newton-Raphson algorithm corresponding to the hyper-reduced model to converge.

4.4.3 B-matrix orthonormal-based solution; indirect version

The obvious question that emerges at this point is the following: if the key ingredient is the perturbation with elements of the orthogonal complement of \mathbb{B}^{IT} (the range of \mathbb{B}^I); why not to *directly* perturb using vectors that belong

¹⁰The main disadvantage of this proposal is the lack of consistency in changing the values of η : the value of the residual is quite problem- and strain history-dependent.

¹¹Demo: Let be $X \in \text{null } \mathbb{B}^{IT}$ and $X' \in \text{range } (\mathbb{B}^I)$. Since $X \in \text{null } \mathbb{B}^{IT}$, then $X^T \mathbb{B}^I = 0$; likewise, since $X' \in \text{range } (\mathbb{B}^I)$, $\exists Y \in \mathbb{R}^{N_u \times 1} \mid \mathbb{B}^I Y = X'$. Hence $X^T \mathbb{B}^I Y = X^T X' = 0$, c.q.d.

¹²It bears mentioning that $\text{rank } (\mathbb{B}^{IT}) = N_u$, otherwise the reduced order model (first stage) could not have been solved. It is not hard to see that $\text{rank } (\mathbb{B}^{IT} \boldsymbol{\Psi}_{\boldsymbol{\eta}}) = N_u$ is also a necessary condition for the Newton-Raphson algorithm corresponding to the hyper-reduced model to converge.

to range (\mathbb{B}^{I^T}) , instead of the certainly circuitous — and more costly — manner of using the residual at the first iteration ? To this end, we perform a SVD over the matrix¹³ \mathbb{B}^I :

$$\mathbb{B}^I = \mathbf{U}_B \mathbf{S}_B \mathbf{V}_B^T. \quad (4.4.15)$$

where, by construction, the matrix of left singular vectors $\mathbf{U}_B \in \mathbb{R}^{N_{gs} \times N_u}$ forms an orthonormal basis for the range of \mathbb{B}^I . Therefore, the i -th column of the snapshot matrix from which the basis Ψ_η are obtained can be computed off-line as:

$$\Sigma_\eta^I{}_i = \Sigma_i^I + \eta \|\Sigma_i^I\| \mathbf{U}_B \alpha_i \quad (4.4.16)$$

where $\alpha_i \in \mathbb{R}^{N_u}$ is a vector of random coefficients ($\alpha_{ij} \in]0, 1[$); the norm of Σ_i^I is introduced to work with a dimensionless perturbing factor.

4.4.3.1 Influence of the perturbing factor

In principle, it seems that the value of the perturbing factor η should affect both the accuracy of the approximation — recall that our ultimate aim is to approximate the stress field from the information contained in a few, strategically selected gauss points — and the robustness of the iterative algorithm. If η is close to zero, the influence of the orthogonal vectors in the snapshot matrix would be negligible and, consequently, the leading (dominant) left singular values — the basis Ψ_η — would be formed by vectors with a negligible component in range (\mathbb{B}^I) , rendering $\text{rank}(\mathbb{B}^{I^T} \Psi_\eta) < N_u$ (necessary condition for convergence). On the other hand, if η is sufficiently large, problems associated with deficient rank of $\mathbb{B}^{I^T} \Psi_\eta$ would disappear, but it seems in all likelihood that large η would have a significant impact on the approximation error. We should, therefore, examine carefully the influence of the perturbing factor in the quality of the approximation. As usual, inspection of the results of a one-dimensional example (in Matlab) will help us to expedite the otherwise laborious task of understanding and unraveling the extent of such an influence.

Suppose that, after performing the *SVD* of the unperturbed snapshot matrix, and by virtue of any truncation criterion, the number of dominant modes is set to N_σ ; on the other hand, let us denote by \check{N}_σ the number of modes of the perturbed snapshot matrix ($\check{N}_\sigma \geq N_\sigma$). First, we study, for several values of \check{N}_σ , and for fixed η , the approximation error made in projecting each column vector of the unperturbed snapshot matrix onto the space spanned by the perturbed basis $\Psi_\eta(\eta)$; i.e.:

$$e_i(\check{N}_\sigma) = \frac{\|(\Psi_\eta \Psi_\eta^T - \mathbf{I}) \Sigma_i^I\|}{\|\Sigma_i^I\|}. \quad (4.4.17)$$

In the results showed in figures 4.1 and 4.2, $\eta = 1$ and $\eta = 100$, respectively, and $N_\sigma = 3$; the additional modes of the perturbed matrix varies from 0 to the number of columns of \mathbf{U}_B , which happens to be equal to 5.

¹³We should note that the SVD is only used for purposed of obtaining an orthonormal basis for the range of the snapshot matrix.

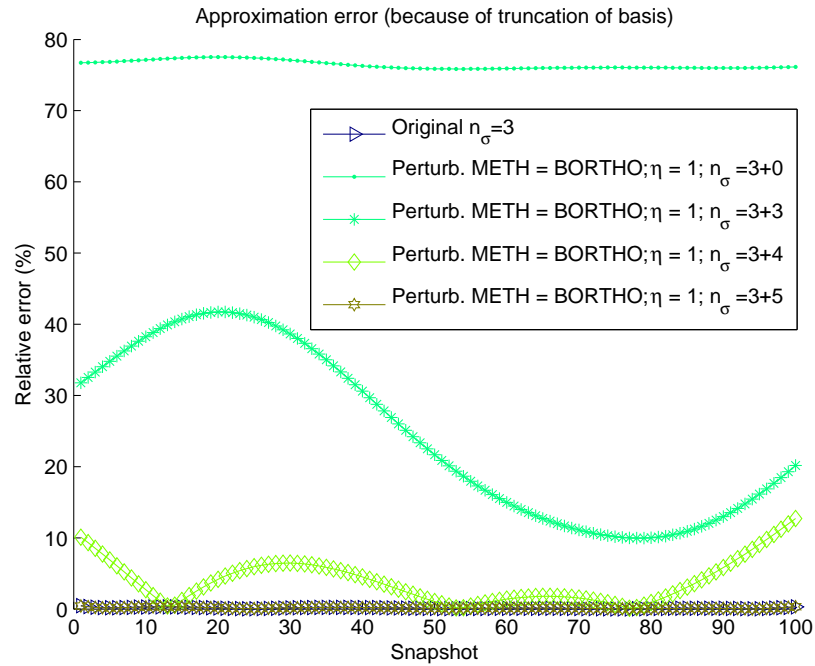


Figure 4.1 Relative error (%) in approximating each column of the snapshot matrix as a linear combination of the perturbed bases; the study is carried for $\eta = 1$ and $\tilde{N}_\sigma = N_\sigma + i$, where $i \in \{0, 3, 4, 5\}$.

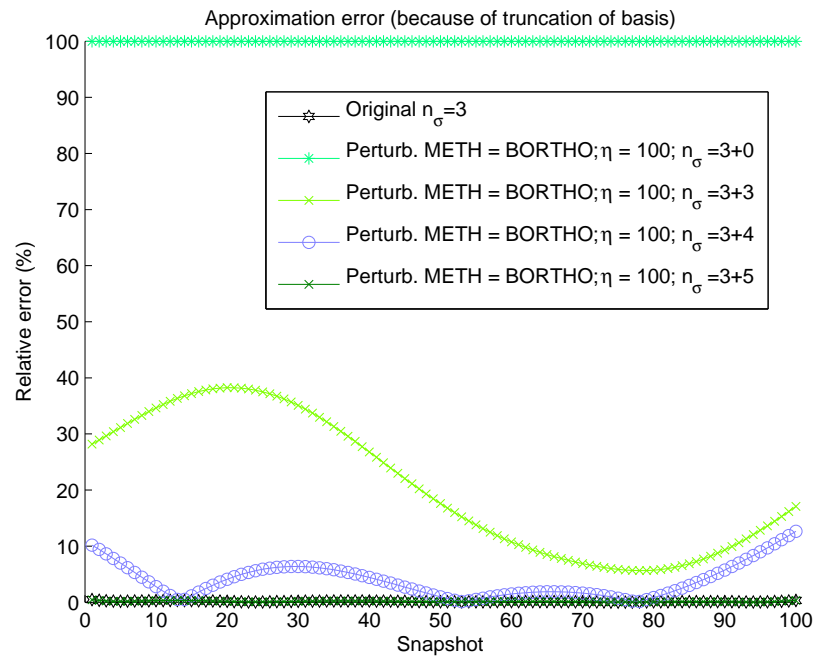


Figure 4.2 Relative error (%) in approximating each column of the snapshot matrix as a linear combination of the perturbed bases. Varying \tilde{N}_σ and $\eta = 100$).

Observation 4.4.1 It can be gleaned from the results displayed in Figs. 4.1 and 4.2 that the approximation error becomes negligible only when $\tilde{N}_\sigma = N_\sigma + N_u$.

To reinforce this tentative conclusion, we show in figure 4.3 the approxima-

tion error for several values of η in the case in which \check{N}_σ is set to $N_\sigma + N_u = 8$. It is clear that the approximation error is not significantly affected by the value of η .

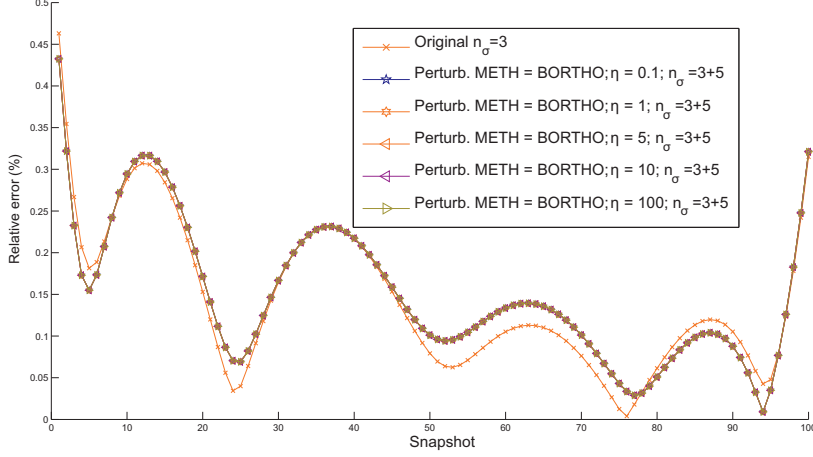


Figure 4.3 Relative error (%) in approximating each column of the snapshot matrix as a linear combination of the perturbed bases. Varying η and $\check{N}_\sigma = N_\sigma + N_u = 8$.

It can be easily appreciated in figure 4.2 that when $\check{N}_\sigma = N_\sigma = 0$, the approximating error is approximately 100% for all snapshots. According to Eq.(4.4.17), this means that $\Psi_\eta \Psi_\eta^T \Sigma_i^I \approx 0$; that is, Σ_i^I and Ψ_η are orthogonal, or equivalently, $\Psi_\eta \in \text{range}(\mathbb{B}^{IT}) = \text{span}\{\mathbf{U}_B\}$. This fact can be easily explained from the relatively large value of η : the dominant modes in this case are but the first N_σ singular left values of the perturbing matrix $\eta \|\Sigma_i^I\| \mathbf{U}_B \alpha_i$.

Observation 4.4.2 *It follows from the above discussion that if η is sufficiently large, the first N_u will be the first singular left values of the perturbing matrix — by construction, the rank of such a matrix is N_u — and the subsequent N_σ modes will be those corresponding to the unperturbed snapshot matrix.*

4.4.3.2 Alternative formulation

The foregoing discussion suggests somehow that, in taking $\check{N}_\sigma = N_\sigma + N_u$, we ensure that the perturbed basis Ψ_η spans a space that contains both the subspace $\text{span}\{\Psi\}$ and $\text{range}(\mathbb{B}^{IT}) = \text{span}\{\mathbf{U}_B\}$ — which, recall, are mutually orthogonal¹⁴. Therefore, any basis vector Ψ_{η_i} can be written as a linear combination of these two set of bases:

¹⁴The validity of such an assertion, however, should be submitted to debate. Strictly, orthogonality is observed if $\mathbb{B}^{IT} \Psi_i = \mathbf{0}$. However, note that Ψ is the basis matrix of the sub-space spanned by stress snapshots that do not meet *exactly* the above condition; rather, each snapshot fulfills $\|\mathbb{B}^{IT} \Sigma\| \leq \text{tol}$, where *tol* is the tolerance specified to determine the termination of the Newton-Raphson solving algorithm. Thus, the higher the convergence tolerance, the more pronounced will be departures from orthogonality. We have examined the observance of orthogonality in a practical case ($r_{snp}^\sigma = 638$). Figure 4.4 contains the graph of the logarithm of the residual $\mathbb{B}^{IT} \Psi_i$ ($i = 1 \cdots r_{snp}^\sigma$). The curve




Figure 4.4 Logarithm of the 2-Norm of the residual $\mathbb{B}^{IT} \Psi_i$ ($i = 1 \dots r_{snp}^\sigma$) (Ψ_i stands for the i -th singular vector arising from the SVD of the matrix corresponding to the training trajectories shown in figure 5.7). The convergence tolerance was set to 10^{-6} .

$$\Psi_{\eta_i} = \Psi \mathbf{d}_i^\sigma + U_B \mathbf{d}_i^u = [\Psi \ U_B] \begin{bmatrix} \mathbf{d}_i^\sigma \\ \mathbf{d}_i^u \end{bmatrix} \quad (4.4.18)$$

for the entire matrix Ψ_η , we have:

$$\Psi_\eta = [\Psi \ U_B] \begin{bmatrix} \mathbf{d}^\sigma \\ \mathbf{d}^u \end{bmatrix} = \Psi_B \mathbf{d}. \quad (4.4.19)$$

In the above $\mathbf{d} = [\mathbf{d}^\sigma \ \mathbf{d}^u]^T \in \mathbb{R}^{\check{N}_\sigma \times \check{N}_\sigma}$ stands for the pertinent coefficient matrix, and $\Psi_B \in \mathbb{R}^{N_g \times \check{N}_\sigma}$ is defined as $\Psi_B = [\Psi \ U_B]$. With the above decomposition at our disposal, the equilibrium equation can be alternatively expressed as:

$$\mathbb{B}^{IT} \Sigma^{II}(\epsilon, \hat{\mathbf{u}}_\mu^{II}) = \mathbb{B}^{IT} \Psi_\eta \mathbf{c}(\epsilon, \hat{\mathbf{u}}_\mu^{II}) = \mathbb{B}^{IT} U_B \mathbf{d}^u \mathbf{c} = \mathbf{0}. \quad (4.4.20)$$

wherein we have exploited the fact that $\mathbb{B}^{IT} \Psi = \mathbf{0}$. Since U_B has complete rank, it follows that the above condition is equivalent to require that¹⁵

$$\mathbf{d}^u \mathbf{c}(\epsilon, \hat{\mathbf{u}}_\mu^{II}) = \mathbf{0}. \quad (4.4.21)$$

Recall that the vector of coefficients \mathbf{c} can be approximated via the gappy data reconstruction method, according to which (see Eq.(4.2.4))

$$\mathbf{c}(\epsilon, \hat{\mathbf{u}}_\mu^{II}) = \Psi_\eta^T \Sigma^{II}(\epsilon, \hat{\mathbf{u}}_\mu^{II}) \approx \left(\hat{\Psi}_\eta^T \hat{\Psi}_\eta \right)^{-1} \hat{\Psi}_\eta^T \hat{\Sigma}^{II}. \quad (4.4.22)$$

exhibits an increasing tendency altered only by moderate oscillations; this means that the “dominant” singular basis vectors are “more orthogonal” to range (\mathbb{B}^{IT}) than the “trailing ones”. A far-reaching (and also interesting) conclusion can be made from this observation: the dominant modes fulfill the equilibrium equation in a more accurate manner than the trailing ones. The fact that the basis vectors with very low singular values can even deteriorate the quality of the solution (they introduce noise).

¹⁵Since $\mathbf{d}^u = U_B^T \Psi_\eta$, this condition is similar to the original one (just replace U_B by \mathbb{B}^{IT})

In the above, $\hat{\Psi}_\eta \in \mathbb{R}^{\hat{N}_g d_s \times \hat{N}_\sigma}$ is a sub-block matrix of Ψ_η constructed by collecting the rows corresponding to \hat{N}_g strategically¹⁶ selected gauss points (whose indices will be denoted by $\hat{\mathcal{I}} = i_1, i_2 \dots i_{\hat{N}_g}$); and $\hat{\Sigma}^{II} \in \mathbb{R}^{\hat{N}_g d_s}$ symbolizes the vector that contains the stresses at such gauss points¹⁷. Defining the boolean matrix $\hat{\mathbb{P}} \in \mathbb{R}^{\hat{N}_g d_s \times N_g d_s}$ as the operator that extracts the rows of Ψ corresponding to the selected gauss points, we can write:

$$\hat{\Psi}_\eta = \hat{\mathbb{P}}\Psi_\eta = [\hat{\mathbb{P}}\Psi \ \hat{\mathbb{P}}U_B] \begin{bmatrix} \mathbf{d}^\sigma \\ \mathbf{d}^u \end{bmatrix} = [\hat{\Psi} \ \hat{U}_B] \begin{bmatrix} \mathbf{d}^\sigma \\ \mathbf{d}^u \end{bmatrix} \quad (4.4.23)$$

and

$$\hat{\Sigma}^{II} = \hat{\mathbb{P}}\Sigma^{II}. \quad (4.4.24)$$

4.4.4 Definitive solution: expansion of the basis matrix

The major criticism that can be leveled at the foregoing “indirect version” is the use of a *random* perturbing matrix. Can we get rid of this admittedly unorthodox manner of expanding the space in which the solution is sought? In particular: Is there any conceptual or mathematical impediment in simply constructing the perturbed bases as $\Psi_\eta = \Psi_B = [\Psi \ U_B]$? In doing so, the second reduction stage would boil down to solving the following problem: for any ϵ , find $\hat{\mathbf{u}}_\mu^{II} \in \mathbb{R}^{N_u}$ such that:

$$\mathbf{c}_U(\epsilon, \hat{\mathbf{u}}_\mu^{II}) = \mathbf{0}. \quad (4.4.25)$$

where $\mathbf{c}_U \in \mathbb{R}^{N_u}$ denotes the last N_u entries of the coefficient vector \mathbf{c} , defined in Eq.(4.4.22). Indeed, any stress vector $\Sigma^I = \Sigma^I(\epsilon, \hat{\mathbf{u}}_\mu^{II})$ can be resolved into two components: one that pertains to span $\{\Psi\}$ and another that resides¹⁸ in span $\{U_B\}$:

$$\Sigma^I(\epsilon, \hat{\mathbf{u}}_\mu^{II}) = \Psi \mathbf{c}_\sigma + U_B \mathbf{c}_U. \quad (4.4.26)$$

A necessary and sufficient condition for $\hat{\mathbf{u}}_\mu^{II}$ to be a solution of the equilibrium problem is that $\Sigma^I(\epsilon, \hat{\mathbf{u}}_\mu^{II}) \in \text{span}\{\Psi\}$. Consequently, the coefficients that multiply the bases U_B must vanish.

Let us now explore more in-depth condition (4.4.25). The coefficient vector \mathbf{c} was given in formula (4.4.22). It was argued there that the approximation contained in Eq.(4.4.25) relies crucially on the invertibility of the following matrix:

$$\hat{M} = \hat{\Psi}_B^T \hat{\Psi}_B = \begin{bmatrix} \hat{M}_{\sigma\sigma} & \hat{M}_{\sigma U} \\ \hat{M}_{\sigma U}^T & \hat{M}_{UU} \end{bmatrix} = \begin{bmatrix} \hat{\Psi}^T \hat{\Psi} & \hat{\Psi}^T \hat{U}_B \\ \hat{U}_B^T \hat{\Psi} & \hat{U}_B^T \hat{U}_B \end{bmatrix}. \quad (4.4.27)$$

Only an appropriate choice of the selection operator¹⁹ $\hat{\mathbb{P}}$ can guarantee the invertibility of the matrix \hat{M} .

¹⁶The choice must ensure the existence of the inverse of $\hat{M} = \hat{\Psi}^T \hat{\Psi}$.

¹⁷It can be easily shown that when $\hat{N}_g = N_g$, the approximation in Eq.(4.4.22) becomes exact.

¹⁸The validity of this assertion rests on the assumption that the solution of the incremental problem is unique. If Σ^I does not represent a self-equilibrated stress field, the component outside the subspace span $\{\Psi\}$ must be necessarily contained in span $\{U_B\}$.

¹⁹Also called the “mask operator”

To extricate the sub-vector \mathbf{c}_U from expression (4.4.22), the inverse of the matrix $\hat{\mathbf{M}}$ is first cast in terms of its block matrices; the final result reads:

$$\begin{aligned} \mathbf{c}_U &= \hat{\mathbf{N}}^{-1} \left(-\hat{\mathbf{M}}_{\sigma U}^T \hat{\mathbf{M}}_{\sigma\sigma}^{-1} \hat{\Psi}^T + \hat{\mathbf{U}}_B^T \right) \hat{\Sigma}^{II} \\ &= \hat{\mathbf{N}}^{-1} \hat{\mathbf{U}}_B^T \left(-\hat{\Psi} (\hat{\Psi}^T \hat{\Psi})^{-1} \hat{\Psi}^T + \mathbf{I} \right) \hat{\Sigma}^{II}. \end{aligned} \quad (4.4.28)$$

where the matrix $\hat{\mathbf{N}}$ is defined as:

$$\hat{\mathbf{N}} = \hat{\mathbf{M}}_{\sigma\sigma} - \hat{\mathbf{M}}_{\sigma U}^T \hat{\mathbf{M}}_{\sigma\sigma}^{-1} \hat{\mathbf{M}}_{\sigma U} = \hat{\mathbf{U}}_B^T \left(\mathbf{I} - \overbrace{\hat{\Psi} (\hat{\Psi}^T \hat{\Psi})^{-1} \hat{\Psi}^T}^{\hat{\mathbf{H}}} \right) \hat{\mathbf{U}}_B. \quad (4.4.29)$$

Inserting the above equation into Eq.(4.4.28), and after trivial manipulation, we finally get

$$\mathbf{c}_U = \left(\hat{\mathbf{U}}_B^T (\mathbf{I} - \hat{\mathbf{H}}) \hat{\mathbf{U}}_B \right)^{-1} \hat{\mathbf{U}}_B^T (\mathbf{I} - \hat{\mathbf{H}}) \hat{\Sigma}^{II}, \quad (4.4.30)$$

where $\hat{\mathbf{H}}$ is a matrix that depends exclusively on the basis $\hat{\Psi}$:

$$\hat{\mathbf{H}} = \hat{\Psi} (\hat{\Psi}^T \hat{\Psi})^{-1} \hat{\Psi}^T. \quad (4.4.31)$$

20

In order to put the problem in an “appealing” format — reminiscent of that presented in Eq.(4.1.5) for the first reduction stage —, we define the *hyper-reduced B-matrix* $\hat{\mathbb{B}}^{II} \in \mathbb{R}^{\hat{N}_g \times N_\sigma + N_u}$ as:

$$\hat{\mathbb{B}}^{II^T} = \left(\hat{\mathbf{U}}_B^T (\mathbf{I} - \hat{\mathbf{H}}) \hat{\mathbf{U}}_B \right)^{-1} \hat{\mathbf{U}}_B^T (\mathbf{I} - \hat{\mathbf{H}}), \quad (4.4.32)$$

so that the hyper-reduced order problem can be posed as: for any ϵ , find $\hat{\mathbf{u}}_\mu^{II} \in \mathbb{R}^{N_u}$ such that:

$$\hat{\mathbb{B}}^{II^T} \hat{\Sigma}^{II}(\epsilon, \hat{\mathbf{u}}_\mu^{II}) = \mathbf{0}. \quad (4.4.33)$$

Observation 4.4.3 *The choice of the basis for the range of \mathbb{B}^I as the left singular vectors arising from the SVD is in principle not relevant; in fact, one can use any basis, in particular, the columns of \mathbb{B}^I itself. In such a case, the expression $\hat{\mathbb{B}}^{II^T}$ adopts the following format:*

$$\hat{\mathbb{B}}^{II^T} = \left(\hat{\mathbb{B}}^{I^T} (\mathbf{I} - \hat{\mathbf{H}}) \hat{\mathbb{B}}^I \right)^{-1} \hat{\mathbb{B}}^{I^T} (\mathbf{I} - \hat{\mathbf{H}}), \quad (4.4.34)$$

where $\hat{\mathbb{B}}^I = \hat{\mathbb{P}} \mathbb{B}^I$.

Observation 4.4.4 *With the choice $\mathbf{U}_B = \mathbb{B}^I$, it is straightforward to demonstrate that when $\hat{N}_g = N_g$, the matrix $\hat{\mathbb{B}}^{II^T}$ degenerates in \mathbb{B}^{II^T} (just note that $\hat{\mathbb{B}}^{II^T} \hat{\mathbf{H}} = \mathbf{0}$ in this case).*

²⁰Should we write $\hat{\mathbb{B}}^{II} \in \mathbb{R}^{\hat{N}_g \times N_\sigma + N_u}$ or $\hat{\mathbb{B}}^{II} \in \mathbb{R}^{\hat{N}_g d \times N_\sigma + N_u}$?

Observation 4.4.5 *The matrix \hat{H} , defined in Eq.(4.4.31), is related with the gappy data reconstruction operator, denoted by \mathcal{R}_σ , through the gauss point selection operator:*

$$\hat{H} = \hat{\mathbb{P}}\mathcal{R}_\sigma = \hat{\mathcal{R}}_\sigma = \overbrace{\hat{\mathbb{P}}\Psi}^{=\hat{\Psi}} \left(\hat{\Psi}^T \hat{\Psi} \right)^{-1} \hat{\Psi}^T. \quad (4.4.35)$$

The operator \mathcal{R}_σ acts on the (self-equilibrated) gappy stress vector $\hat{\Sigma}^{II}$ to give the full-gauss point vector Σ^{II} :

$$\Sigma^{II} = \mathcal{R}_\sigma \hat{\Sigma}^{II}. \quad (4.4.36)$$

Using the identity $\hat{H} = \hat{\mathcal{R}}_\sigma$, we can rephrase Eq.(4.4.34) as:

$$\hat{\mathbb{B}}^{II^T} = \left(\hat{\mathbb{B}}^{IT} (\mathbf{I} - \hat{\mathcal{R}}_\sigma) \hat{\mathbb{B}}^I \right)^{-1} \hat{\mathbb{B}}^{IT} (\mathbf{I} - \hat{\mathcal{R}}_\sigma). \quad (4.4.37)$$

4.5 Analysis of approximation errors

4.5.1 Sampling and reconstruction errors

Discrepancies between stress responses computed by the full-order and hyper-reduced-order models can be decomposed in the following manner:

$$\begin{aligned} \Sigma^h - \Sigma^{II} &= \overbrace{\Sigma^h - \Sigma^I}^{e_\Sigma^I} + \overbrace{\Sigma^I - \Sigma^{II}}^{e_\Sigma^{II}} \\ &= e_\Sigma^I + e_\Sigma^{II}. \end{aligned} \quad (4.5.1)$$

Here, the term e_Σ^I represents²¹ the error incurred in the projection stage — discussed previously in section 3.3.3.1 —, and e_Σ^{II} encompasses the errors due exclusively to the *reconstruction* of the stress field from the values computed at a few, strategically selected gauss points. The component e_Σ^{II} can be further split into *sampling* ($e_\Sigma^{II,samp}$) and *reconstruction* ($e_\Sigma^{II,rec}$) components:

$$e_\Sigma^{II} = e_\Sigma^{II,rec} + e_\Sigma^{II,samp}. \quad (4.5.2)$$

The sampling error gives an indication of the richness of information contained in the stress snapshot matrix; it is defined formally as

$$e_\Sigma^{II,samp} := \Sigma^I - \Sigma^{II}(r_{snp}^\sigma, N_g), \quad (4.5.3)$$

where $\Sigma^{II}(r_{snp}^\sigma, N_g)$ denotes the stress solution provided by the HROM in the limit of no truncation, that is, when $N_\sigma = r_{snp}^\sigma$ (r_{snp}^σ is the rank of the snapshot matrix), and when the entire set of gauss points is used for reconstruction purposes ($\hat{N}_g = N_g$). The remaining error is the part due to reconstruction:

$$e_\Sigma^{II,rec} := \Sigma^{II}(r_{snp}^\sigma, N_g) - \Sigma^{II}, \quad (4.5.4)$$

and it depends on the truncation level (N_σ), on the one hand, and on the number (\hat{N}_g) and location²² of the selected points.

²¹We should note that e_Σ^I and e_Σ^{II} are in general not orthogonal; thus, it is possible that either $\|\Sigma^h - \Sigma^{II}\| \leq \|e_\Sigma^{II}\|$ or $\|\Sigma^h - \Sigma^{II}\| \leq \|e_\Sigma^I\|$

²²Thus, the reconstruction error depends also on the quality of the procedure used for select such locations. A (cursory) review of existing selection algorithms was given in section 1.2.2.1

4.5.2 Consistency

To discuss the issue of consistency, it proves conceptually convenient to summarize first the micro-cell equilibrium equations corresponding to each level of reduction; this is done in the box below.

$$\mathbb{B}^{hT} \Sigma(\mathbf{u}_\mu^h, \boldsymbol{\epsilon}_t) = \mathbf{0}, \quad \text{with } \mathbb{B}^{hT}: \mathbb{R}^{N_g d_s} \rightarrow \mathbb{R}^{n d}; \quad (4.5.5)$$

$$\mathbb{B}^{IT} \Sigma(\mathbf{u}_\mu^I, \boldsymbol{\epsilon}_t) = \mathbf{0}, \quad \text{with } \mathbb{B}^{IT}: \mathbb{R}^{N_g d_s} \rightarrow \mathbb{R}^{N_u}; \quad (4.5.6)$$

$$\hat{\mathbb{B}}^{II^T} \hat{\Sigma}(\hat{\mathbf{u}}_\mu^{II}, \boldsymbol{\epsilon}_t) = \mathbf{0}, \quad \text{with } \hat{\mathbb{B}}^{II^T}: \mathbb{R}^{\hat{N}_g d_s} \rightarrow \mathbb{R}^{N_u}. \quad (4.5.7)$$

where

$$\mathbb{B}^{IT} = \Phi^T \mathbb{B}^{hT}, \quad (4.5.8)$$

$$\hat{\mathbb{B}}^{II^T} = \mathbb{B}^{IT} \mathcal{R}_B, \quad (4.5.9)$$

$$\mathcal{R}_B = \Psi_B \left(\hat{\Psi}_B^T \hat{\Psi}_B \right)^{-1} \hat{\Psi}_B^T \quad (4.5.10)$$

$$\Psi_B = [\Psi U_B], \quad \hat{\Psi}_B = \hat{\mathbb{P}} \Psi_B, \quad (4.5.11)$$

$$\hat{N}_g \geq N_u + N_\sigma. \quad (4.5.12)$$

$\Phi \in \mathbb{R}^{n d \times N_u} \rightarrow$ Micro-disp. fluct. POD basis matrix

$\Psi \in \mathbb{R}^{N_g d_s \times N_\sigma} \rightarrow$ Micro-stress. POD basis matrix

$U_B \in \mathbb{R}^{N_g d_s \times N_u} \rightarrow$ Orthogonal basis matrix for range (\mathbb{B}^I)

$\hat{\mathbb{P}} \in \mathbb{R}^{\hat{N}_g d_s \times N_g d_s} \rightarrow$ (Boolean) mask operator .

Box 4.5.1 *Micro-cell equilibrium equations (in condensed, matrix format) corresponding to the full-order model (Eq.(4.5.5)), the reduced-order model (Eq.(4.5.6)) and the hyperreduced-order model (Eq.(4.5.7)). The process of reduction in dimension and complexity is reflected in the progressive ‘‘compression’’ of the B-matrix operator: $\mathbb{B}^{hT} \Rightarrow \mathbb{B}^{IT} \Rightarrow \hat{\mathbb{B}}^{II^T}$ (recall that the whole reduction process is predicated on the assumption that $N_u \ll n$ and $\hat{N}_g, N_\sigma \ll N_g$).*

As discussed earlier (section 3.3.3.2), consistency demands that when the prescribed strain history pertains to the set of training trajectories ($\boldsymbol{\epsilon}_t \in \boldsymbol{\epsilon}_t^{snp}$), and in the limit of no data compression ($N_u = r_{snp}^u$, $N_\sigma = r_{snp}^\sigma$ and $\hat{N}_g = N_g$), both the solution of the full-order problem represented by Eq.(4.5.5) and the solution of the hyper-reduced order problem stated in Eq.(4.5.7) coincide ($\hat{\mathbf{u}}_\mu^{II} = \mathbf{u}_\mu^h$). Consistency between the FOM and ROM solutions was already proven in section 3.3.3.2; thus, it only remains to demonstrate that $\hat{\mathbf{u}}_\mu^{II} = \mathbf{u}_\mu^I$ in the limiting case of no truncation of the stress basis and when the entire set of gauss points is used.

To prove this, first note that, in the limiting case of no data compression, the mask operator becomes the identity and, as a consequence, the reconstruction operator \mathcal{R}_B defined in Eq.(4.5.10) degenerates into the or-

thogonal projector onto $\text{span}\{\Psi_B\}$, that is:

$$\mathcal{R}_B = \Psi_B (\Psi_B^T \Psi_B)^{-1} \Psi_B^T = \Psi_B \Psi_B^T. \quad (4.5.13)$$

Accordingly, the HROM equilibrium equation (4.5.7) can be rephrased as:

$$\mathbb{B}^{IT} \Psi_B \Psi_B^T \Sigma(\hat{\mathbf{u}}_\mu^{II}, \epsilon_t) = 0. \quad (4.5.14)$$

On the other hand, since $\epsilon_t \in \epsilon_t^{snp}$, we know that $\Sigma(\mathbf{u}_\mu^I, \epsilon_t) \in \text{span}\{\Psi\}$; therefore, we can say that

$$\Sigma(\mathbf{u}_\mu^I, \epsilon_t) = \Psi \Psi^T \Sigma(\mathbf{u}_\mu^I, \epsilon_t). \quad (4.5.15)$$

Furthermore, since U_B is a basis of the range of \mathbb{B}^I , it follows that $U_B^T \Sigma(\mathbf{u}_\mu^I, \epsilon_t) = \mathbf{0}$. In view of this observation, and by virtue of Eq.(4.5.15), it can be thus asserted that

$$\Sigma(\mathbf{u}_\mu^I, \epsilon_t) = \Psi_B \Psi_B^T \Sigma(\mathbf{u}_\mu^I, \epsilon_t). \quad (4.5.16)$$

Inserting the above equation into the ROM equilibrium equation (4.5.6), we finally arrive at

$$\mathbb{B}^{IT} \Sigma(\mathbf{u}_\mu^I, \epsilon_t) = \mathbb{B}^{IT} \Psi_B \Psi_B^T \Sigma(\mathbf{u}_\mu^I, \epsilon_t) = \mathbf{0}. \quad (4.5.17)$$

Comparison of the above expression with Eq.(4.5.14) permits to conclude that²³ the solutions provided by the hyper-reduced order model and the reduced-order model are identical in the limit of no data compression.

4.6 Computation of macroscopic stresses

The ultimate aim of solving the equilibrium problem associated to the fine scale is not the determination of the microscopic fluctuation field *per se*, but rather the computation of the macroscopic stress tensor via *homogenization* — i.e., volume averaging of the microscopic stress distribution:

$$\boldsymbol{\sigma}(\epsilon) = \frac{1}{V_\mu} \int_{\Omega_\mu^s} \boldsymbol{\sigma}_\mu d\Omega, \quad (4.6.1)$$

where V_μ stands for the total volume of the RVE. Approximating the integral by gauss quadrature, we get:

$$\boldsymbol{\sigma}(\epsilon) = \frac{1}{V_\mu} \sum_{g=1}^{N_g} w_g \boldsymbol{\sigma}_{\mu_g}(\epsilon). \quad (4.6.2)$$

By virtue of Eq.(4.4.36), the micro-stress field $\boldsymbol{\sigma}_\mu$ can be, in turn, approximated as:

$$\Sigma(\epsilon) \approx \mathcal{R}_\sigma \hat{\Sigma}(\epsilon) \rightarrow \begin{bmatrix} \boldsymbol{\sigma}_{\mu 1} \\ \boldsymbol{\sigma}_{\mu 2} \\ \vdots \\ \boldsymbol{\sigma}_{\mu N_g} \end{bmatrix} \approx \begin{bmatrix} \mathcal{R}_{\sigma 1} \\ \mathcal{R}_{\sigma 2} \\ \vdots \\ \mathcal{R}_{\sigma N_g} \end{bmatrix} \hat{\Sigma}, \quad (4.6.3)$$

²³Needless to say, this holds strictly only if the full-order model is well-posed (uniqueness of solution).

that is,

$$\boldsymbol{\sigma}_{\mu_g}(\boldsymbol{\epsilon}) \approx \mathcal{R}_{\sigma_g} \hat{\boldsymbol{\Sigma}}(\boldsymbol{\epsilon}) \quad g = 1, 2 \cdots N_g, \quad (4.6.4)$$

where \mathcal{R}_{σ_g} stands for the sub-block matrix formed by the rows of the reconstruction operator (see Eq.(4.4.36)) corresponding to the g -th gauss point. Replacement of $\boldsymbol{\sigma}_\mu$ in Eq.(4.6.2) by the above coefficient-function approximation finally yields:

$$\boldsymbol{\sigma}(\boldsymbol{\epsilon}) = \mathcal{T} \hat{\boldsymbol{\Sigma}}(\boldsymbol{\epsilon}), \quad (4.6.5)$$

where the operator $\mathcal{T} \in \mathbb{R}^{d_s \times \hat{N}_g d_s}$ is defined as follows:

$$\mathcal{T} := \frac{1}{V_\mu} \sum_{g=1}^{N_g} w_g \mathcal{R}_{\sigma_g} \quad (4.6.6)$$

Note that \mathcal{T} can be entirely *computed offline*, since it does not depend on the value of the macrostrain $\boldsymbol{\epsilon}$. Further insight is gained by rephrasing expression (4.6.5) as

$$\boldsymbol{\sigma}(\boldsymbol{\epsilon}) = [\mathcal{T}_1 \mathcal{T}_2 \cdots \mathcal{T}_{\hat{N}_g}] \begin{bmatrix} \boldsymbol{\sigma}_\mu(\mathbf{y}_{p_1}, \boldsymbol{\epsilon}) \\ \boldsymbol{\sigma}_\mu(\mathbf{y}_{p_2}, \boldsymbol{\epsilon}) \\ \vdots \\ \boldsymbol{\sigma}_\mu(\mathbf{y}_{p_{\hat{N}_g}}, \boldsymbol{\epsilon}) \end{bmatrix} = \sum_{g=1}^{\hat{N}_g} \mathcal{T}_g \boldsymbol{\sigma}_\mu(\mathbf{y}_{p_g}, \boldsymbol{\epsilon}), \quad (4.6.7)$$

where $\boldsymbol{\sigma}_{\mu_{p_g}}(\boldsymbol{\epsilon}) = \boldsymbol{\sigma}_\mu(\mathbf{y}_{p_g}, \boldsymbol{\epsilon})$ denotes the value of the micro-stress vector at the g -th sample point \mathbf{y}_{p_g} . According to this expression, the computation of the macroscopic stress involves only determination of the micro-stresses at the selected quadrature points — by solving the hyperreduced-order equilibrium equation (4.5.7) — and multiplication of such stresses by their corresponding (pre-computed) *weighting matrices* $\mathcal{T}_g \in \mathbb{R}^{d_s \times d_s}$ ($g = 1, 2 \cdots \hat{N}_g$). It is interesting to note that, in fact, there is no need to *reconstruct* the micro-stress field through direct application of the reconstruction operator \mathcal{R}_σ and, as a consequence, the *operation count of the online stage depends exclusively on the dimension (\hat{N}_g) of the hyperreduced-order model*.

Observation 4.6.1 *Equation (4.6.7) represents the approximation of the integral appearing in Eq.(4.6.1) using the POD basis computed from the stress snapshot matrices; accordingly, the operators $\mathcal{T}_g \in \mathbb{R}^{d_s \times d_s}$ ($g = 1, 2 \cdots \hat{N}_g$) can be interpreted, in this approximation, as the weights associated to the values of the integrand at locations $\mathbf{y}_{p_1}, \mathbf{y}_{p_2}, \cdots, \mathbf{y}_{p_{\hat{N}_g}}$. The peculiarity here is that such weights are not scalars, as in classical quadrature schemes, but matrices — hence the appellation *weighting matrices*.*

4.6.1 Homogenized algorithmic tangent moduli

The proposed hyperreduced-order strategy affords computational savings not only in the update of the macroscopic stress updates, but also in the computation of the macroscopic *algorithmic* tangent moduli, a task which proves rather costly in homogenization procedures. In the following, we present — without dwelling unduly on theoretical considerations — the derivation of this tensor in both the full and hyperreduced-order models; for a more detailed and formal account of the former, the reader is referred to de Souza Neto and Feijóo (2006).

4.6.1.1 Full-order model

The algorithmic tangent moduli, denoted by \mathbf{C} , is obtained by applying the directional derivative relative to a *macroscopic* strain increment $\Delta\boldsymbol{\epsilon}$ to expression²⁴ (4.6.1):

$$\begin{aligned}
\mathcal{D}\boldsymbol{\sigma}[\Delta\boldsymbol{\epsilon}] &= \mathbf{C}\Delta\boldsymbol{\epsilon} = \frac{1}{V_\mu} \int_{\Omega_\mu^s} \mathcal{D}\boldsymbol{\sigma}_\mu(\overbrace{\boldsymbol{\epsilon} + \nabla^s \mathbf{u}_\mu}^{\boldsymbol{\epsilon}_\mu})[\Delta\boldsymbol{\epsilon}] d\Omega \\
&= \frac{1}{V_\mu} \int_{\Omega_\mu^s} \overbrace{\frac{\partial \boldsymbol{\sigma}_\mu}{\partial \boldsymbol{\epsilon}_\mu}}^{\mathbf{C}_\mu} \mathcal{D}\boldsymbol{\epsilon}_\mu[\Delta\boldsymbol{\epsilon}] d\Omega \\
&= \overbrace{\frac{1}{V_\mu} \int_{\Omega_\mu^s} \mathbf{C}_\mu d\Omega}^{\mathbf{C}^{taylor}} \Delta\boldsymbol{\epsilon} + \frac{1}{V_\mu} \int_{\Omega_\mu^s} \mathbf{C}_\mu \mathcal{D}\nabla^s \mathbf{u}_\mu[\Delta\boldsymbol{\epsilon}] d\Omega \\
&= \mathbf{C}^{taylor} \Delta\boldsymbol{\epsilon} + \frac{1}{V_\mu} \int_{\Omega_\mu^s} \mathbf{C}_\mu \nabla^s \mathcal{D}\mathbf{u}_\mu[\Delta\boldsymbol{\epsilon}] d\Omega.
\end{aligned} \tag{4.6.8}$$

Expressing \mathbf{u}_μ as a linear combination of the finite element basis functions (see Eq.(3.2.1)), we get:

$$\begin{aligned}
\mathcal{D}\boldsymbol{\sigma}[\Delta\boldsymbol{\epsilon}] &= \mathbf{C}^{taylor} \Delta\boldsymbol{\epsilon} + \frac{1}{V_\mu} \int_{\Omega_\mu^s} \mathbf{C}_\mu \mathbf{B} d\Omega \overbrace{\mathcal{D}\hat{\mathbf{u}}_\mu[\Delta\boldsymbol{\epsilon}]}{=\mathcal{P}_{u_\mu} \Delta\boldsymbol{\epsilon}} \\
&= \left(\mathbf{C}^{taylor} + \overbrace{\frac{1}{V_\mu} \int_{\Omega_\mu^s} \mathbf{C}_\mu \mathbf{B} d\Omega}^{\mathbf{C}^{fluc}} \mathcal{P}_{u_\mu} \right) \Delta\boldsymbol{\epsilon}.
\end{aligned} \tag{4.6.9}$$

The homogenized algorithmic tangent tensor \mathbf{C} , thus, consists in the sum of two linear operators. One of them, denoted in the foregoing equation as²⁵ \mathbf{C}^{taylor} , is simply the volume average of the algorithmic tangent tensor consistent with the microscopic constitutive law. The other contribution, symbolized by \mathbf{C}^{fluc} , reflects the fact that the displacement fluctuations themselves depend on the history of macro-strains. The tangential operator

²⁴To keep the notation uncluttered, we have dropped the superindex “h” (finite element solution) — it should read, for instance, $\mathcal{D}\boldsymbol{\sigma}^h[\Delta\boldsymbol{\epsilon}]$. Moreover, tensors have been expressed directly in Voigt’s notation.

²⁵This term corresponds to the homogenized tangent moduli of a micro-cell in which the fluctuations are assumed to be zero, and, thus, $\boldsymbol{\epsilon}_\mu = \boldsymbol{\epsilon}$. Under this kinematical assumption, known by the name of *Taylor’s assumption* —hence the superscript — or *rule of mixtures*, the homogenization problem admits a trivial solution; however, it should be noted that such a solution indiscriminately ignores possible interactions between the distinct solid micro-constituents and the effects induced by micro-porosity (presence of voids). For instance, the macro-stress predicted by Taylor’s model in the case of a porous cell with a single solid phase reads $\boldsymbol{\sigma} = \frac{V_\mu^s}{V_\mu} \boldsymbol{\sigma}_\mu$ (V_μ^s is the volume of the solid phase, excluding the voids). This solution, thus, does not account for the influence of the void distribution on the macro-scale properties; if the micro-constituent obeys an isotropic constitutive equation, the homogenization process will preserve this isotropy regardless of how voids are distributed. Needless to say, this is patently false — just imagine a cell containing a single, asymmetrically placed void.

\mathcal{P}_{u_μ} , that relates nodal fluctuations and macro-strains, is obtained by applying the directional derivative $\mathcal{D}(\bullet)[\Delta\epsilon]$ to the RVE equilibrium equation (3.2.6):

$$\begin{aligned}
& \int_{\Omega_\mu^s} \mathbf{B}^T \mathcal{D} \sigma_\mu[\Delta\epsilon] d\Omega = \\
& = \int_{\Omega_\mu^s} \mathbf{B}^T \mathbf{C}_\mu d\Omega \Delta\epsilon + \int_{\Omega_\mu^s} \mathbf{B}^T \mathbf{C}_\mu \nabla^s \mathcal{D} \mathbf{u}_\mu[\Delta\epsilon] d\Omega \\
& = \int_{\Omega_\mu^s} \mathbf{B}^T \mathbf{C}_\mu d\Omega \Delta\epsilon + \overbrace{\int_{\Omega_\mu^s} \mathbf{B}^T \mathbf{C}_\mu \mathbf{B} d\Omega}^{\mathbf{K}_\mu} \mathcal{D} \hat{\mathbf{u}}_\mu[\Delta\epsilon] = \mathbf{0}.
\end{aligned} \tag{4.6.10}$$

The desired tangential operator \mathcal{P}_{u_μ} emerges from the above by simply solving for $\mathcal{D} \hat{\mathbf{u}}_\mu[\Delta\epsilon]$ — in the manner outlined in section (3.2.1).

4.6.1.2 Hyperreduced-order model

The derivation of the homogenized tangent operator consistent with the hyperreduced-order model follows the same line as above. Application of the directional derivative $\mathcal{D}(\bullet)[\Delta\epsilon]$ to expression (4.6.7) yields:

$$\begin{aligned}
\mathcal{D} \sigma[\Delta\epsilon] &= \sum_{g=1}^{\hat{N}_g} \mathcal{T}_g \mathcal{D} \sigma_{\mu_{p_g}}[\Delta\epsilon] \\
&= \underbrace{\sum_{g=1}^{\hat{N}_g} \mathcal{T}_g \mathbf{C}_{\mu_{p_g}}}_{\hat{\mathbf{C}}^{taylor}} \Delta\epsilon + \sum_{g=1}^{\hat{N}_g} \mathcal{T}_g \mathbf{C}_{\mu_{p_g}} \mathcal{D} \nabla^s \mathbf{u}_\mu[\Delta\epsilon].
\end{aligned} \tag{4.6.11}$$

Expressing \mathbf{u}_μ as a linear combination of the reduced-order basis functions, we get:

$$\begin{aligned}
\mathcal{D} \sigma[\Delta\epsilon] &= \hat{\mathbf{C}}^{taylor} \Delta\epsilon + \underbrace{\left(\sum_{g=1}^{\hat{N}_g} \mathcal{T}_g \mathbf{C}_{\mu_{p_g}} \mathbf{B}_{p_g}^I \right)}_{\hat{\mathbf{C}}^{fluc}} \overbrace{\mathcal{D} \hat{\mathbf{u}}_\mu^{II}[\Delta\epsilon]}^{=\mathcal{P}_{u_\mu} \Delta\epsilon} \\
&= (\hat{\mathbf{C}}^{taylor} + \underbrace{\left(\sum_{g=1}^{\hat{N}_g} \mathcal{T}_g \mathbf{C}_{\mu_{p_g}} \mathbf{B}_{p_g}^I \right)}_{\hat{\mathbf{C}}^{fluc}} \mathcal{P}_{u_\mu}) \Delta\epsilon.
\end{aligned} \tag{4.6.12}$$

To obtain a closed-form expression for the operator \mathcal{P}_{u_μ} , the directional derivative $\mathcal{D}(\bullet)[\Delta\epsilon]$ is applied to the hyperreduced-order version of the equilibrium equation (Eq.(4.5.7)):

$$\begin{aligned} \hat{\mathbb{B}}^{II^T} \hat{\Sigma} &= \sum_{g=1}^{\hat{N}_g} \hat{\mathbb{B}}_g^{II^T} \sigma_{\mu p_g} = \mathbf{0} \xrightarrow{\mathcal{D}(\bullet)[\Delta \epsilon]} \\ & \sum_{g=1}^{\hat{N}_g} \hat{\mathbb{B}}_g^{II^T} \mathbf{C}_{\mu p_g} \Delta \epsilon + \sum_{g=1}^{\hat{N}_g} \hat{\mathbb{B}}_g^{II^T} \mathbf{C}_{\mu p_g} \mathbf{B}^I_{p_g} \mathcal{D} \hat{u}_\mu^{II}[\Delta \epsilon] = 0. \end{aligned} \quad (4.6.13)$$

Solving for $\mathcal{D} \hat{u}_\mu^{II}[\Delta \epsilon]$, we finally arrive at the desired expression:

$$\mathcal{P}_{u_\mu} = - \left(\sum_{g=1}^{\hat{N}_g} \hat{\mathbb{B}}_g^{II^T} \mathbf{C}_{\mu p_g} \mathbf{B}^I_{p_g} \right)^{-1} \sum_{g=1}^{\hat{N}_g} \hat{\mathbb{B}}_g^{II^T} \mathbf{C}_{\mu p_g} \quad (4.6.14)$$

Chapter 5

Assessment of approximation errors: a case study

5.1 Introduction

The aim of this chapter is twofold. First, to analyze the errors incurred in approximating the finite element or full-order solution (FOM) by the successive reduced-order approximations (ROM and HROM) for particular training and testing macro-strain trajectories; and second, to evaluate the actual efficiency of the proposed model reduction strategy (in terms of speedup). To this end, we shall study the behavior, under plane strain conditions, of a RVE containing a single centered circular hole (see figure 5.1), and governed by a damage model endowed with linear hardening ($\sigma_y = 1000 \text{ MPa}$, $H = 0.1$)




Figure 5.1 Geometry of the micro-cell (1024 linear triangle elements; 574 nodes).

5.2 First reduction stage (ROM)

In the following, the notions of truncation and sampling errors discussed in section 3.3.3.1 are enlivened by two examples. The three macroscopic strain histories used for generating the snapshots in the first of these examples are represented in figure 5.2. Observe that, in each of this strain histories, one of the strain components follows a piecewise linear, ascending-descending path — with a maximum of $\epsilon_C = 2 \cdot 10^{-3}$, ($C = x, y, xy$) —, while the magnitude of the other two components is set to zero. Besides, the time domain for each strain history is discretized into 220 equally spaced steps. In figure 5.3, on the other hand, we show the macroscopic strain trajectory employed for testing the quality of the approximation.

Observation 5.2.1 *The choice of the three strain histories shown in figure 5.2 is not based in any rigorous, statistically sound sampling procedure, but rather is dictated by physical intuition. The task of “training”¹ the micro-cell is somehow akin to the experimental process whereby material parameters of standard phenomenological models are calibrated in a laboratory. In this analogy, the micro-cell plays the role of the corresponding experimental specimen, and the macro-strain training trajectories represents the loading paths of the pertinent calibration tests. Specifically, the strain paths displayed in figure 5.2 correspond to: a “tensile test” in the direction of the x -axis (snapshots 1-220); a “tensile test” in the direction of the y -axis (snapshots 221-440); and a “shear test” (snapshots 441-660). As opposed to the situation encountered in laboratory experiments, however, in the training process of the micro-cell one has privileged information regarding the underlying micro-structure, and this may aid to restrict the number of tests necessary to characterize the response. For instance, the behavior of the micro-cell under study is governed by a rate-independent, symmetric, damage model; thus, we know beforehand that including compression tests would prove redundant (because of the symmetry of the constitutive relationship); likewise, due to the rate-independence character of the problem, it is not necessary to study the response under varying rates of deformation.*

Figure 5.4 contains the graph of the error estimate defined as:

$$e_{\sigma}^{\max}(N_u) = \frac{\max_{n=1,2,\dots} \|\Sigma^h(t_n) - \Sigma^I(t_n, N_u)\|_2}{\sqrt{N_g d_s} \sigma_y} 100, \quad (5.2.1)$$

i.e., the maximum of the 2-norm of the difference between the FOM and ROM stress solutions corresponding to the trajectories shown in figure 5.3, and for varying levels of truncation (ranging from $N_u = 9$ to $N_u = r_{snp}^u = 290$). The error falls abruptly at approximately $N_u = 10$, remaining thereafter practically flat. In the limit of no truncation ($N_u = r_{snp}^u = 292$), the discrepancy between FOM and ROM solutions is below 2% of the ultimate strength σ_y — this is an estimation of the *sampling error* for the testing trajectory under consideration. Observe also that the curve decays in a monotonic manner; this indicates that, although the basis are constructed in an optimal fashion only for the micro-displacement fluctuations (the state

¹The term training, which, incidentally, is borrowed from the neural network literature, is used throughout the text to refer to the offline generation of snapshots.




Figure 5.2 Macro-strain history used for generating the displacement snapshots.




Figure 5.3 Macro-strain history used for testing the quality of the reduced-order approximation.

variable), the optimality requirement alluded to in section 3.3.3.2 is also observed in terms of stresses.

5.2.0.3 Analysis with different testing trajectories

Next, we examine the error incurred in approximating the evolution of stresses using, for testing purposes, a different strain history, which is displayed in figure 5.5. The paths that follow each component of the macroscopic strain tensor lie entirely in the *negative half-plane* —unlike the training trajectories (figure 5.2), that take only positive values. Furthermore, the strain paths in figure 5.5 reach a maximum absolute value that is *one order of magnitude greater* than the peak value attained by the training




Figure 5.4 Error estimation defined in Eq.(5.2.1) ($e_{\sigma}^{max}(N_u)$) versus the level of truncation N_u . Error measures have been divided by the ultimate strength of the material ($\sigma_y = 1000 \text{ MPa}$) to get dimensionless values.

data. This strain history, thus, poses a more severe test of the ability of the reduced-order model to approximate the full-order response.




Figure 5.5 Macro-strain history used for examining the quality of the approximation (“demanding” situation).

Figure 5.6 contains the plot of the error measure versus the level of truncation for the trajectories shown in figures 5.3 and 5.5 (labeled as *ST1* and *ST2* in figure 5.6). As expected, the degree of discrepancy at low level of truncations for the *ST2* case is substantially larger than for the trajectory

labeled as ST1 (60% against 8% at $N_u = 5$). However, note that the drift between these curve decreases notably with increasing number of modes. At $N_u = 20$ the error estimate is below 10%, a value that can be deemed satisfactory from an engineering point of view; in the limit of no truncation (sampling error), the level of discrepancy is only of 6%.

The above analysis seems to suggest that even in situations in which parts of the testing macro-strain trajectory lie far out of the range of the training macro-strain functions, the reduced-order model is able to provide reasonably accurate estimations using only a few modes ($N_u = 20$)— it proves its ability to make accurate *extrapolations*, in other words².




Figure 5.6 Maximum of the 2-norm error versus level of truncation for the testing trajectories shown in 5.3 (strain history 1) and 5.5 (strain history 2). Error measures have been divided by the ultimate strength ($\sigma_y = 1000 \text{ MPa}$) to get dimensionless values.

5.2.0.4 Analysis with different training trajectories

In this section, we are concerned with the assessment of the error incurred in approximating the stress evolution induced by a given testing trajectory when different set of training data are used. The testing strain path chosen for the analysis is the one shown previously in figure 5.5. Two sets of training data, on the other hand, have been used for comparison purposes: the three strain trajectories shown previously (figure 5.2), and the strain histories displayed in figure 5.7. Observe that the latter set includes also combined tensile-shear loading/unloading paths, and, besides, the peak loads are greater (by an order of magnitude) than those in figure 5.2; hence the training data contained in figure 5.7 constitutes a more comprehensive sample of the parametric space than the pure shear and pure tensile paths shown in figure 5.2. Another perceptible difference is that the the descending (unloading) branches of the paths depicted in figure 5.7 do not reach the zero level . This has been done in anticipation of the fact that, being

²We should mention in this respect that one of the acclaimed advantages of POD/-Galerkin reduced-order models is precisely their ability to make “physically-based extrapolations” (as opposed to “black-box“ approaches, e.g., artificial neural networks, which are somehow “agnostic“ to the underlying physics.)

the unloading process elastic, the span of the matrix formed by the snapshots corresponding to this range cannot be greater than one, and therefore, collecting too many snapshots is not necessary³.

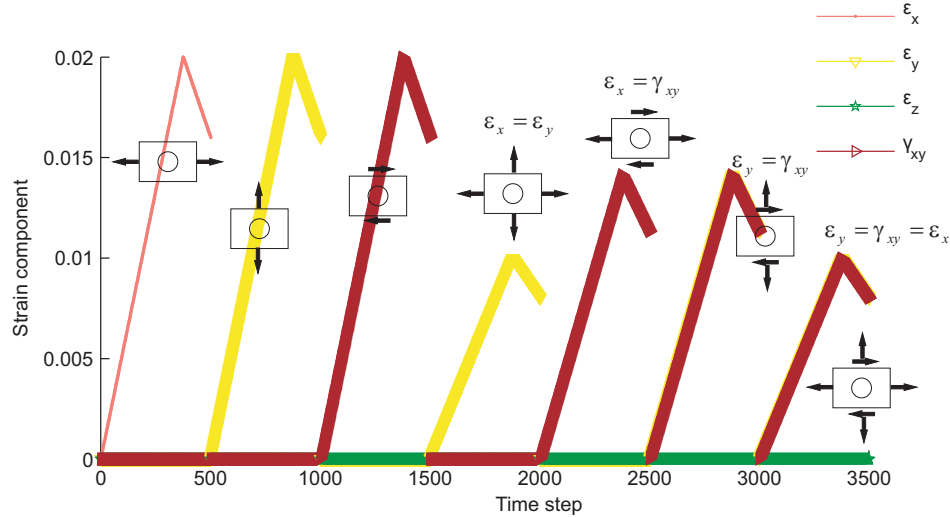


Figure 5.7 Macro-strain histories used for generating displacement fluctuations snapshots. These data represent a more comprehensive sample of the parametric space than the strain paths depicted in figure 5.2.

Figure 5.8 contains the graphs of the 2-norm error stress versus the level of truncation for these two distinct sets of training strain histories. Results clearly concurs with expectations: the solution provided by the reduced-order model that uses the training data shown in figure 5.7 represents a significantly more accurate approximation than the one that uses the training trajectories contained in figure 5.2; with only 3 modes of the former we get the same accuracy that with 20 modes of the latter. This comparison conveys the relevance of properly selecting the macro-strain training histories if a parsimonious representation of the micro-cell is to be obtained.

To shed more light into the — admittedly — diffuse topic of “sampling” of the parametric space, we plot in figure 5.9 the 2-norm of the *sampling* error incurred at each time step accompanied by the graph of the global Newton-Raphson iterations required at each time step to achieve equilibrium in the micro-cell. Inspection of this figure indicates that the error peak is located for both cases in the time interval between the 10 and 15 steps. Interestingly, yet not surprisingly, this corresponds to the time interval at

³In principle, just two snapshots would suffice to characterize the displacement response in this elastic range, and, thus, one might be naively tempted to include *only* two snapshots in the snapshot ensemble for purposes of POD basis construction. However, recall that the POD is nothing but a multidimensional data fitting intended to obtain a sequence of orthogonal vectors whose spans best approximate the range of the snapshot matrix. If the elastic response is poorly represented in the snapshot ensemble, the POD would regard as unimportant the contribution of these snapshots, and, as a consequence, the dominant modes would hardly contain any information of this range; one would be forced to take a large number of modes to accurately replicate the apparently trivial elastic behavior. This observation highlights the importance of assuring a balanced presence in the snapshot ensemble of solutions corresponding to relevant, phenomenological distinct behaviors. As suggested in appendix A.4, a circuitous way of achieving this is to affect the snapshots by certain, aptly chosen weighting coefficients — in the spirit of the snapshot-weighting scheme developed by (Carlberg and Farhat, 2008).




Figure 5.8 Maximum of the 2-norm error versus level of truncation for the testing trajectory contained in 5.5 when using two distinct sets of training strain histories: the one displayed in figure 5.2 (labeled as “TRD 1”) and the one shown in figure 5.7 (labeled as “TRD 2”). Error measures have been divided by the ultimate strength ($\sigma_y = 1000 \text{ MPa}$) to get dimensionless values.




Figure 5.9 (a) 2-Norm of the sampling stress error at each time step for the testing trajectory shown in 5.5. Error measures have been divided by the ultimate strength ($\sigma_y = 1000 \text{ MPa}$) to get dimensionless values. (b) Number of iterations performed until reaching convergence at each step (full-order model)

which more iterations are required to achieve equilibrium. Further insight is gained by analyzing the propagation of damage through the cell. In figure 5.10, we plot the portion of finite elements (in %) affected by damage as a function of time. Observe that the interval at issue corresponds to the period during which damage spreads out through the cell; i.e., when both elements behaving inelastically and elements that remains in the elastic range coexist. It is apparent than, to reduce the sampling error, the time stepping in this interval should be refined and more solution snapshots corresponding to this range included in the snapshot matrix.

5.3 Second reduction stage (HROM)

5.3.1 Optimal choice of truncation levels and number of sample points

As explained in section 4.5, the quality of the hyper-reduced-order approximation is controlled by three integer variables: the level of truncation of the




Figure 5.10 (a) Portion of elements affected by damage at each time step for the testing trajectory shown in 5.5. (b) Number of iterations required to achieve equilibrium.

micro-displacement basis (N_u), the level of truncation of the micro-stresses basis (N_σ), and the number of gauss points (\hat{N}_g) at which to evaluate — through the corresponding micro-constitutive equations — the micro-stress response. The only constraint in the choice of this triplet is the (necessary) condition for the existence of the reconstruction operator \mathcal{R}_B , that reads

$$\hat{N}_g \geq N_\sigma + N_u. \quad (5.3.1)$$

Since the computational cost of increasing the level of truncations is not significant — the actual burden is the evaluation of the stress response at the \hat{N}_g selected points —, and considering that the quality of the approximation is, obviously, invariably improved with the expansion of the reduced-order spaces, it follows that, for a given number of points \hat{N}_g , the level of truncation should be as large as possible⁴; accordingly, the most favorable situation is encountered when:

$$\hat{N}_g = N_\sigma + N_u. \quad (5.3.2)$$

Incidentally, note that this limiting case represents the closest situation to a genuine interpolatory problem.

It only remains, thus, to ascertain which combination of N_σ and N_u , for a given \hat{N}_g , is most advantageous in terms of accuracy. In order to assist in the answering of such a question, we have performed — using the training and testing trajectories shown in figures 5.7 and 5.3, respectively — two battery of tests; in these tests, \hat{N}_g is kept constant and equal to 20 and 30, while N_σ and N_u vary so that $\hat{N}_g = N_\sigma + N_u$. Figure 5.11 contains the graph, for each of these tests, of the error estimate defined as:

$$e_\sigma^{max}(N_u) = \frac{\max_{n=1,2,\dots} \|\Sigma^h(t_n) - \Sigma^{II}(t_n, N_u, N_\sigma, \hat{N}_g)\|_2}{\sqrt{N_g d_s} \sigma_y} 100, \quad (5.3.3)$$

⁴This argument is also supported by the one-dimensional analysis carried out in appendix C.

versus the ratio between the number of stress modes and the number of sample points. Observe that the minimum is located for both curves at around $N_\sigma/\hat{N}_g = 0.5$; thus, according to this numerical experiment, the *optimum relation between levels of truncations is* $N_\sigma = N_u$ (equal levels of truncation for both displacements and stresses).




Figure 5.11 Error measure defined in Eq.(5.3.3) versus the quotient between the number of stress modes N_σ and the number of sample points $\hat{N}_g = N_\sigma + N_u$ for two different values of \hat{N}_g .

A plausible, somehow intuitive explanation for this finding goes as follows. The quality of the stress solution depends largely, through the pertinent constitutive relationships, on the quality of the displacement solution, and vice versa. Therefore, it seems pointless to increase the level of truncation of the stress basis without accompanying such an increase with a pertinent enlargement of the displacement solution space. On the other hand, the plot of the *a priori*⁵ truncation error estimator defined as:

$$e_b^{trun} = \frac{R - M}{R} \sqrt{\frac{\sum_{i=M+1}^R \lambda_i^2}{\sum_{i=1}^R \lambda_i^2}}. \quad (5.3.4)$$

versus⁶ the level of truncation M ($M = N_\sigma$ for stresses and $M = N_u$ for displacements), see figure 5.12, indicates that stress and displacement truncation errors are in the same order of magnitude and decrease at similar rates⁷. In view of these facts, it is apparent that there is no reason to take N_u appreciably larger than N_σ , or vice-versa; the optimum relation, thus,

⁵The “a priori” qualifier refers to the fact that this estimation can be obtained in the *offline* stage, as it only depends of the outcome of the SVD of the snapshot matrices; see appendix C.2.2 for further details.

⁶here, λ_i and R denote the i -th singular value and rank of the corresponding snapshot matrix.

⁷Observe also that both curves exhibit similar patterns, having inflection points at the exact same locations. It may be interesting to investigate this peculiarity in future

should be close to $N_u = N_\sigma$, as the numerical experiment carried out above suggested.

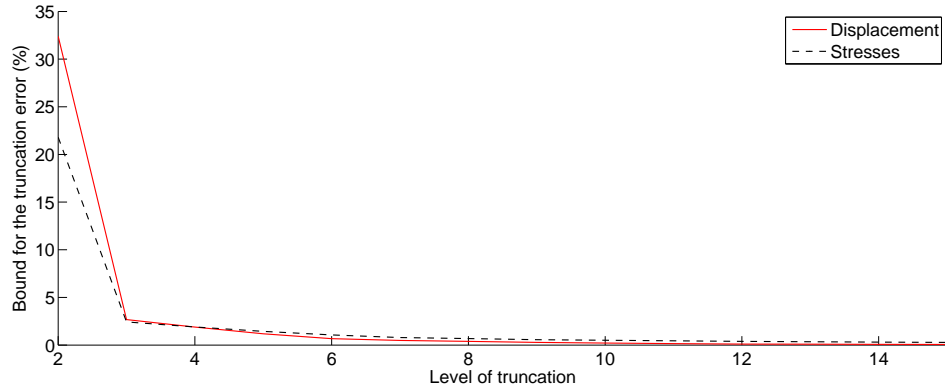


Figure 5.12 Bound for the displacement and stress truncation error defined in Eq.(5.3.4) versus the level of truncation.

5.3.2 Influence of the greedy sampling strategy

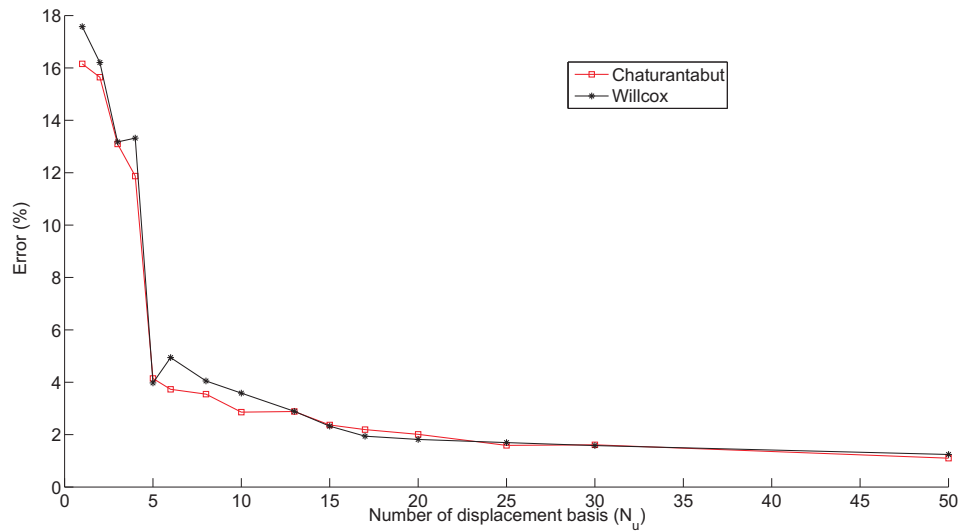


Figure 5.13 Comparison between greedy algorithms (Willcox (2006) and Chaturantabut and Sorensen (2010)). Frobenius norm of the HROM stress errors versus the number of displacement modes N_u (with respect to σ_y), with $N_\sigma = N_u$ and $\hat{N}_g = 2N_u$.

The approximation error depends not only on the number of *gappy* points used for reconstructing the stress field, but also on the location of such points within the micro-cell domain and, therefore, on the particular algorithm employed for determining such locations. In order to preserve the continuity of the presentation, the in-depth discussion of the various selection methods proposed in the literature and the influence on the error response of the precise placement of the *gappy* points have been deferred to appendix C. Here

developments and see if the resemblance between truncation error curves is also observed with other, more complex micro-models constitutive laws.

we limit ourselves to compare the error responses associated to the greedy sampling procedures proposed by Willcox (2006), on the one hand, which is based on the minimization of the condition number of the *restricted* basis matrix, and the one pioneered by Barrault et al. (2004) and later extended to discrete settings by Chaturantabut and Sorensen (2010), on the other hand — the so-called Discrete Empirical Interpolation Method (DEIM). Figure 5.13 displays the plot of the error estimator defined in Eq.(5.3.3) versus the number of displacement modes N_u for these two sampling procedures. Inspection of this figure permits to conclude that differences in performance between these algorithms are not significant. Although for moderate levels of truncation the plot associated with the method advocated by Chaturantabut and Sorensen (2010) provides slightly more accurate results and exhibit a less erratic behavior than the other one, these apparent advantages becomes imperceptible as N_u increases. We should point out that this conclusion is in total agreement with the remarks made in appendix C in the context of approximation of (discontinuous) one-dimensional functions using POD basis.

5.3.3 Gains in computational efficiency

We come now to the main concern of the present work — its *raison d'être* — which is the issue of *computational efficiency*. The upcoming analysis is devoted to roughly⁸ evaluate the performance gains that, for a given error level, can be achieved by using the hyperreduced-order model in the computation of the response of the fine scale.

5.3.3.1 Local speedup

Let us begin by defining the *local speedup factor* associated to the computation of the response of a single micro-cell:

$$s_\mu := \frac{t_{FOM}(N_g)}{t_{HROM}(\hat{N}_g)}, \quad (5.3.5)$$

where t_{FOM} and t_{HROM} denote the times required to compute, respectively, the full- and hyperreduced-order macro-stress responses induced by a given macro-strain history. Figure 5.14.(a) contains the plot of this local speedup factor versus the number of sample points \hat{N}_g for the macro-strain history employed in section 5.2 (see figure 5.3). For purposes of discussing the trade-off between speedup and quality of the low-rank approximation, this plot is accompanied by the graph of the dimensionless error measure defined as

$$e_{macro}^{max}(\hat{N}_g) = \frac{\max_{n=1,2,\dots} \|\boldsymbol{\sigma}^h(t_n) - \boldsymbol{\sigma}^{II}(t_n, \hat{N}_g)\|_\infty}{\sigma_y} 100, \quad (5.3.6)$$

⁸A rigorous and unbiased — i.e., not contingent upon computer architecture, programming language and/or coding style features — assessment of the benefits that accrue from using the proposed low-rank approximation should be carried out by measuring the number of computer operations (*flops*) involved in each relevant step of the algorithm; the work by DeVore et al. (2001) provides thorough guidelines to accomplish this task. In the present work, however, we do not pursue such a degree of rigor, but rather limit ourselves to simply determine and compare, using an in-house Matlab program, cpu elapsed times for the computation of the full-order and hyperreduced-order fine scale macro-stress responses induced by a representative strain history.

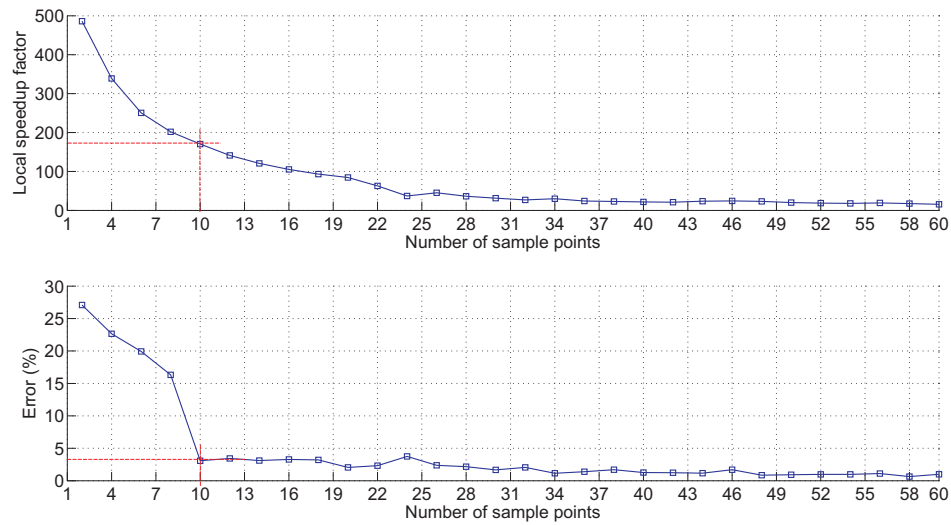


Figure 5.14 (a) Local speedup factor versus the number of sample points for the computation of the response corresponding to the macro-strain history employed in section figure 5.3. (b) Maximum difference between full- and hyperreduced-order macro-stress responses — divided by the ultimate strength of the material, see Eq.(5.3.6) — versus the number of sample points.

i.e., the maximum difference between computed full- and hyperreduced-order *macroscopic* stress⁹ responses. Note that the speedup at $\hat{N}_g = 2$ is slightly below 500, and at $\hat{N}_g = 22$ is around 50; dividing the total number of points $N_g = 1024$ by these figures we get $1024/500 \approx 500$ and $1024/22 \approx 50$. It follows, then, that the speedup decreases with the ratio N_g/\hat{N}_g at an approximately linear rate¹⁰. The error curve displayed in the accompanying figure, on the other hand, exhibits a notable drop ($\Delta e \approx 25\%$) from $\hat{N}_g = 2$ to $\hat{N}_g = 10$, remaining thereafter practically flat¹¹ — the decay of error from $\hat{N}_g = 10$ to $\hat{N}_g = 30$ is only of 3%, while the speedup decreases tenfold. Therefore, it becomes apparent that a satisfactory compromise between accuracy and computational cost can be achieved by taking $\hat{N}_g = 10$; we have in this case only 4% of error with a speedup of two order of magnitudes over the full-order model. For completeness, we show in figure 5.15 the evolution for this case of the components of the macroscopic stresses calculated with both FOM and HRM. On the other hand, the location of the $\hat{N}_g = 10$ sample points used in the computations — selected by the algorithm proposed¹² Chaturantabut and Sorensen (2010) — is displayed in

⁹Recall that the ultimate aim of computing the fine scale is the calculation of the macroscopic stress tensor σ via homogenization; thus, the error estimate Eq.(5.3.6) provides a more meaningful — from a practical point of view — indication of the quality of the approximation than the error measure defined previously in Eq.(5.3.3). The latter is more suited for examining convergence properties of the hyperreduced approximation — the minimization problem that underlies the SVD is posed in terms of the 2-norm,.

¹⁰This is a manifestation of the fact that the calculation of the stresses dominates the total cost of the computation of the fine scale response.

¹¹Notice that the decreasing pattern of this curve mimic that of the plot of the a priori error bound displayed in figure 5.12.

¹²Recall that the *expanded* basis matrix is formed by the first N_σ stress modes and a set of orthogonal basis vectors for the range of \mathbb{B}^I ($[\Psi \mathbf{U}_B]$). Since greedy algorithms, by definition, are hierarchical, the exact location of the first $N_\sigma = 5$ selected points depends exclusively on the stress basis Ψ . Incidentally, it can be appreciated in figure 5.16.(a)

figure 5.16.

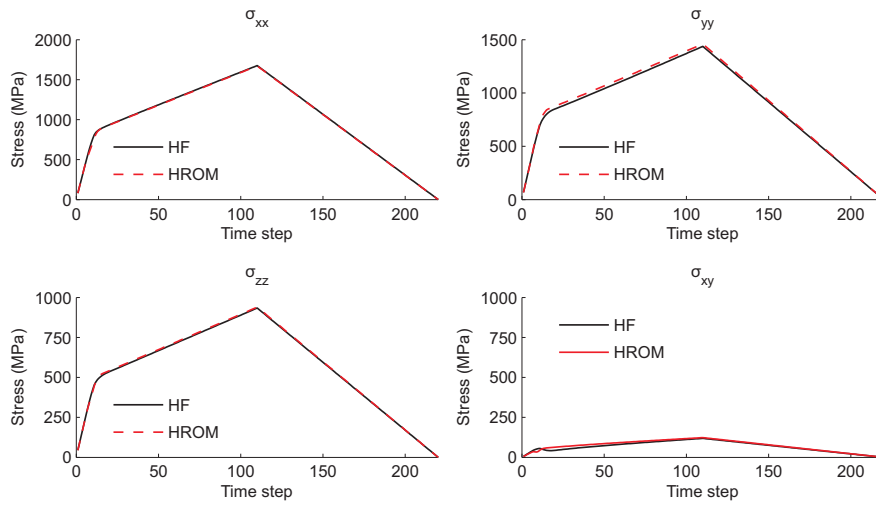


Figure 5.15




Figure 5.16 Location of the $\hat{N}_g = 10$ sample points selected by the greedy algorithm proposed by Chaturantabut and Sorensen (2010) (DEIM). (a) Points associated to the first $N_\sigma = 5$ stress modes. (b) Points associated to the orthogonal basis for the range of \mathbb{B}^I (of rank $N_u = 5$).

Observation 5.3.1 *The foregoing discussion considers only the trade-off between speedup and error levels for varying \hat{N}_g and fixed N_g . But what about the behavior when N_g varies and \hat{N}_g is kept constant? How would affect an increase of N_g to the total error? For instance, if the mesh shown above is uniformly refined by, say, a factor of 3, and the number of sample points is kept to $\hat{N}_g = 10$, will the error level be the same (4%)? Or should one*

that four of the such points are distributed along the perimeter of the void, where stress concentrations are likely to take place. This is somehow consistent with the observations made at the end of appendix C when discussing the connection between the shape of the spatial modes and the location of optimal sample points.

increase also the number of sample points by a factor of 3 so as to preserve accuracy? These questions are thoroughly addressed in appendix B. It is theoretically and numerically proven therein that the truncation error — the major contributor to the total error — is relatively insensitive to mesh refinement and, as a consequence, the speedups provided by the hyperreduced-order model will invariably grow — at a rate roughly proportional to the factor of refinement — as the spatial discretization of the RVE becomes finer.

5.3.3.2 Global speedup

A fair discussion of the gains in efficiency provided by the hyperreduced-order model should take into account, aside from the local speedup, the additional cost associated to the offline stage. Let us denote by N_g^{macro} and T^{macro} the number of quadrature points of the finite element mesh at the coarse level, and the number of time steps used for the solution of the problem, respectively. The overall efficiency afforded by the proposed model reduction strategy can be loosely estimated by the following *global speedup factor*:

$$s_{macro} := \frac{t_{FOM}}{t_{off} + t_{HROM}} = \frac{t_0 + N_g^{macro} T^{macro} \Delta t_{\mu}^{FOM}}{t_{off} + t_0 + N_g^{macro} T^{macro} \Delta t_{\mu}^{FOM}}. \quad (5.3.7)$$

Here t_{off} represents the cost of the offline stage; Δt_{μ}^{FOM} and Δt_{μ}^{HROM} stand for the time required to update the macro-stresses using the RVE full- and hyperreduced-order models, respectively; and t_0 denotes the computational overhead associated to other relevant operations (inversion of global stiffness matrix, for instance). The total offline cost t_{off} is mainly due to the generation of the snapshots — the construction of global basis and selection of sample points are comparatively negligible — and, therefore, we can write $t_{off} \approx T^{train} \Delta t_{\mu}^{FOM}$, with T^{train} being the number of snapshots computed in the training process. Likewise, the ratio between Δt_{μ}^{FOM} and Δt_{μ}^{HROM} is approximately equal to the speedup factor defined earlier ($\Delta t_{\mu}^{FOM} / \Delta t_{\mu}^{HROM} \approx s_{\mu}$). These considerations enable Eq.(5.3.7) to be rephrased as

$$s_{macro} = \frac{t_0 + N_g^{macro} T^{macro} s_{\mu} \Delta t_{\mu}^{HROM}}{s_{\mu} T^{train} \Delta t_{\mu}^{HROM} + t_0 + N_g^{macro} T^{macro} \Delta t_{\mu}^{HROM}}. \quad (5.3.8)$$

According to the above expression, model reduction achieves superior runtime performance ($s_{macro} > 1$) only if

$$N_g^{macro} T^{macro} > \frac{s_{\mu}}{s_{\mu} - 1} T^{train}. \quad (5.3.9)$$

If the total cost is dominated by the update of macrostresses — a fairly reasonable assumption —, the term t_0 can be ignored in expression 5.3.8, and it reduces to

$$s_{macro} = r \frac{s_{\mu}}{s_{\mu} + r}. \quad (5.3.10)$$

The constant r in the above is defined as

$$r = \frac{N_g^{macro} T^{macro}}{T^{train}}. \quad (5.3.11)$$

The plot of this admittedly crude estimation of the global speedup factor versus r is shown in figure 5.17 for the local speedup obtained in the previous section ($s_\mu = 150$). It can be gleaned from this plot that global speedups of above two order of magnitudes are achieved in this case at $r > 500$. If, for instance, $T^{train} = 3500$ (as in the training trajectories depicted in section 5.2.0.4, figure 5.7), the simulation of a problem with, say, $T^{macro} = 1000$ time steps and $N_g^{macro} = 2000$ quadrature points at the macroscale level will generously repaid back the offline investment. Obviously, the benefits of model reduction becomes more evident when the same training output is reused in several problems.




Figure 5.17 Global speedup (Eq.(5.3.10)) as a function of $r = N_g^{macro} T^{macro} / T^{train}$

It almost goes without saying that the proposed model reduction strategy affords dramatic savings not only in terms of CPU time, but also in terms of memory requirement. Indeed, in full-order simulations, history data has to be updated at a number of integration points equal to the number of Gauss points in the macro problem multiplied by the number of Gauss points in the unit cell ($N_g^{macro} N_g$); thus, memory savings will be roughly of the order of the ratio \hat{N}_g / N_g — in the analyzed case $10/1024 \approx 1\%$.

Appendix A

The Singular Value Decomposition (SVD)

A.1 Introduction

Scientific literature abounds with references to the *Singular value decomposition* (SVD) and the *Proper Orthogonal decomposition* (POD); it is applied to many problems, not only within engineering circles, but also to other disparately distinct disciplines. For instance, the SVD is exploited to design learning machines with the ability of acquiring human-like knowledge from the same sources as humans (see Landauer et al. (1998)). The history of the SVD method is reviewed in Stewart (1993) and Kerschen et al. (2005); in the latter, the authors emphasize the sheer number of distinct denominations that receives this technique: “... it is emphasized that the method appears in various guises in the literature and is known by other names depending on the area of application, namely PCA in the statistical literature, empirical orthogonal function in oceanography and meteorology, and factor analysis in psychology and economics.”. Likewise, Kerschen et al. (2005) discuss the dubious physical interpretation of the POD modes. Limitations of SVD are discussed in a literarily exquisite manner in Smith (2002); the connection of SVD with principal component analysis (PCA) are also disclosed therein.

A.2 Formulation

An illuminating account on the *singular value decomposition* (SVD) technique is given in Press et al. (1993): “SVD methods are based on the following theorem of linear algebra [...]: Any M by N matrix \mathbf{A} whose number of rows¹ M is greater than or equal to its number of columns N , can be written as the product of an M by N column-orthogonal matrix \mathbf{U} , an N by N diagonal matrix \mathbf{W} with positive or zero elements (the singular values), and the transpose of an N by N orthogonal matrix \mathbf{V} :

$$\mathbf{A} = \mathbf{U} \cdot \mathbf{W} \cdot \mathbf{V}^T. \quad (\text{A.2.1})$$

It is instructive to note that (page 53) “ SVD explicitly constructs orthonormal bases for the *nullspace and range of a matrix*. Specifically,

¹The SVD decomposition can also be carried out when $M < N$ (see page 52 of Press et al. (1993))

the columns of \mathbf{U} whose same-numbered elements w_j are nonzero are an orthonormal set of basis vectors that span the range; the columns of \mathbf{V} whose same-numbered elements w_j are zero are an orthonormal basis for the nullspace. [...] When the vector \mathbf{b} ($\mathbf{A} \cdot \mathbf{x} = \mathbf{b}$) on the right-hand side is not zero, the important question is whether it lies in the range of \mathbf{A} or not. If it does, then the singular set of equations does have a solution \mathbf{x} ; in fact it has more than one solution, since any vector in the nullspace (any column of \mathbf{V} with a corresponding zero w_j) can be added to \mathbf{x} in any linear combination.” (Press et al., 1993).

There is a close connection between singular value decomposition and the linear least-squares problem; this is didactically shown in page 54 of Press et al. (1993). The connection between SVD and classical eigenvalue problems is also clearly explained in Press et al. (1993), as well as the role of the SVD in the approximation of matrices (page 58).

A.3 Correlation between snapshots: the covariance matrix

When the snapshot matrix is normalized ($\mathbf{X}_{\sigma i}^I = \mathbf{\Sigma}^I / \|\mathbf{\Sigma}^I\|$), the “covariance” matrix² defined (in components) as

$$\mathbf{S}_{\sigma ij}^I = \mathbf{X}_{\sigma i}^{IT} \mathbf{X}_{\sigma j}^I, \quad -1 \geq \mathbf{S}_{\sigma ij}^I \leq 1, \quad (\text{A.3.1})$$

furnishes information about the degree of linear correlation between snapshots. If $|\mathbf{S}_{\sigma ij}^I| \approx 1$, it can be stated that the i -th and j -th snapshots are highly correlated, whereas if $|\mathbf{S}_{\sigma ij}^I| \approx 0$, it is said that the i -th and j -th snapshots are virtually uncorrelated (linearly independent vectors).


It might prove interesting to investigate the magnitude of the entries of the covariance matrix for the case under study (see figure 5.2). Since it is nearly impossible to numerically display such a huge matrix (660x660), we prefer to represent it in a more colorful fashion (see figure A.1). Red entries means that $|\mathbf{S}_{\sigma ij}^I| \approx 1$ (highly correlated) whereas blue connotes that the snapshots are virtually uncorrelated $|\mathbf{S}_{\sigma ij}^I| \approx 0$. It can be seen from figure A.1 that the snapshot corresponding to tensile strain histories are barely correlated with the snapshots associated with shear strain histories (see figure A.1(a)); on the other hand, the degree of correlation within the same strain history snapshots is markedly high (see figure A.1(b)).

A.4 Some properties of the SVD

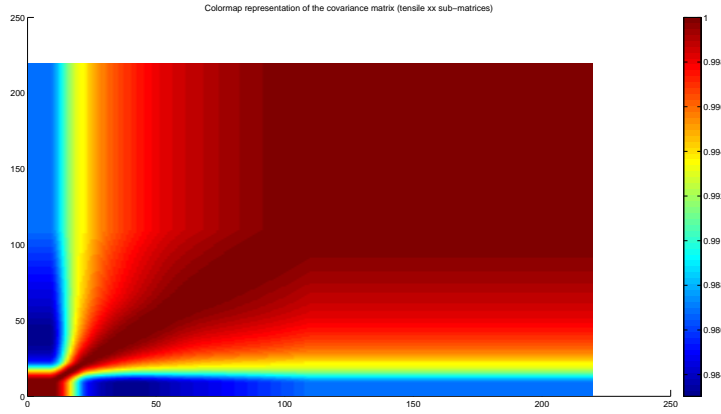
The following self-evident³ properties might prove of invaluable help in devising strategies for aptly sampling the parametric space — selection of solution snapshots — and in discussing issues related to the sensitivity of truncation error to mesh refinements.

²As pointed out in Vaseghi (2007), the SVD is essentially an eigenanalysis of the matrix $\mathbf{X}_{\sigma}^{IT} \mathbf{X}_{\sigma}^I$, which, incidentally, can be interpreted as a covariance matrix. The SVD simply extracts dominant features by disposing of those orthogonal dimensions or features that have insignificant variance

³By virtue of the uniqueness of the singular value decomposition.



(a) Entire matrix



(b) Sub-matrix corresponding to the first 220 snapshots (tensile strain history)

Figure A.1 Colormap representation of the normalized covariance matrix.

Theorem A.4.1. Let Σ_1, Σ_2 be two mutually orthogonal vectors in \mathbb{R}^N . The SVD of the matrix defined as

$$\mathbf{X} = \underbrace{[\Sigma_1 \ \Sigma_1 \ \Sigma_1 \ \cdots]}_{M_1 \text{ times}} \underbrace{[\Sigma_2 \ \Sigma_2 \ \Sigma_2 \ \cdots]}_{M_2 \text{ times}}, \text{ with } M_1 > M_2, \quad (\text{A.4.1})$$

is given by

$$\mathbf{X} = [\Sigma_1 \ \Sigma_2] \begin{bmatrix} \sqrt{M_1} & 0 \\ 0 & \sqrt{M_2} \end{bmatrix} \begin{bmatrix} \frac{\alpha_1^T}{\sqrt{M_1}} & \mathbf{0} \\ \mathbf{0} & \frac{\alpha_2^T}{\sqrt{M_2}} \end{bmatrix}. \quad (\text{A.4.2})$$

where $\alpha_i \in \mathbb{R}^{M_i}$, $\alpha_i = [1 \ 1 \ \cdots]^T$ ($i = 1, 2$).

Theorem A.4.2. Let $\mathbf{X} = \mathbf{U}\mathbf{S}\mathbf{V}^T$ be the SVD of a matrix $\mathbf{X} \in \mathbb{R}^{N \times S}$. The SVD of the matrix $\bar{\mathbf{X}} \in \mathbb{R}^{N \times M \times S}$ defined as:

$$\bar{\mathbf{X}}^T = \underbrace{[\mathbf{X}^T \ \mathbf{X}^T \ \mathbf{X}^T \ \cdots]}_{M \text{ times}}, \quad (\text{A.4.3})$$

is given by $\bar{\mathbf{X}} = \bar{\mathbf{U}}\bar{\mathbf{S}}\bar{\mathbf{V}}^T$, where

$$\bar{\mathbf{U}}^T = \frac{1}{\sqrt{M}} \overbrace{[\mathbf{U}^T \mathbf{U}^T \mathbf{U}^T \dots]}^{M \text{ times}}, \quad (\text{A.4.4})$$

$$\bar{\mathbf{V}}^T = \mathbf{V}^T, \quad (\text{A.4.5})$$

and

$$\bar{\mathbf{S}} = \sqrt{M}\mathbf{S}. \quad (\text{A.4.6})$$

Appendix B

Influence of mesh refinement on the truncation error

B.1 Introduction

It was pointed out in section 5.3.3 that the speedup at the fine scale level s_μ afforded by the hyperreduced-order model depends essentially on the ratio between the number of quadrature points of the finite element mesh N_g and the number of sample points \hat{N}_g :

$$s_\mu \sim \frac{N_g}{\hat{N}_g}. \quad (\text{B.1.1})$$

The *discretization* error e^{fem} — i.e., the error between the continuum and the finite element solutions — is controlled by N_g , whereas \hat{N}_g governs the truncation error e^{trun} between this finite element solution and the hyperreduced-order response. Since the online cost is entirely independent of \hat{N}_g , one may be tempted, in the quest for higher accuracy and realism of the overall representation, to increase N_g — that is, to refine the finite element mesh of the micro-cell. The obvious, yet not trivial, question that emerges in doing so is: should this increase of N_g be accompanied also by an increase of the number of sample points \hat{N}_g ? E.g., if we double the number of quadrature points of the full-order model, should we also double the number of sample points in order to maintain the same truncation error?

Let $\mathbf{X} = \mathbf{U}\mathbf{S}\mathbf{V}^T$ and $\bar{\mathbf{X}} = \bar{\mathbf{U}}\bar{\mathbf{S}}\bar{\mathbf{V}}^T$ be the SVD of the stress snapshot matrix associated with two different meshes containing, respectively, N_g and αN_g ($\alpha > 1$) quadrature points. As discussed in section 5.3.1 the truncation error is bounded above by a quantity that depends only on the magnitude of the singular values:

$$e_b^{trun} = \frac{R-M}{R} \sqrt{\frac{\sum_{i=M+1}^R \lambda_i^2}{\sum_{i=1}^R \lambda_i^2}}. \quad (\text{B.1.2})$$

where M and R stands for the level of truncation and the rank of the matrix, respectively (recall that $M = \hat{N}_g/2$). It follows from this expression that the

influence of refinement of the discretization on the truncation error can be assessed by merely analyzing and comparing the singular values $\lambda_i = \mathbf{S}_{ii}$ and $\bar{\lambda}_i = \bar{\mathbf{S}}_{ii}$ of both \mathbf{X} and $\bar{\mathbf{X}}$, respectively. However, in a general situation, this comparison between the singular values before and after mesh refinement can be only performed numerically. As usual in error assessment in finite element contexts, theoretical results can be only established in the limiting case of the mesh size tending to zero.

B.2 Asymptotic analysis

In the limiting case of $N_g \rightarrow \infty$, for each point of the refined mesh there is at least one (neighboring) point in the old mesh with essentially the same properties. As a consequence, the snapshot matrix associated to the finer grid can be legitimately written as

$$\bar{\mathbf{X}} = \mathbf{A}\mathbf{X}, \quad (\text{B.2.1})$$

with $\mathbf{A} \in \mathbb{R}^{N_g \alpha \times N_g}$ being the corresponding boolean matrix¹. In the special case of uniform mesh refinement² and $\alpha \in \mathbb{N}$, the snapshot matrix corresponding to the refined mesh is directly expressible as:

$$\bar{\mathbf{X}}^T = \overbrace{[\mathbf{X}^T \ \mathbf{X}^T \ \mathbf{X}^T \ \dots]}^{\alpha \text{ times}}. \quad (\text{B.2.2})$$

By virtue of the result derived in Eq.(A.4.6), we know that the matrix of singular values of at snapshot matrix $\bar{\mathbf{X}}$ formed by α stacked copies of \mathbf{X} is given by:

$$\bar{\mathbf{S}} = \sqrt{\alpha}\mathbf{S}. \quad (\text{B.2.3})$$

It becomes apparent from the preceding expression that, when $N_g \rightarrow \infty$, the *truncation error remains unaltered upon (uniform) mesh refinements*.

B.3 Numerical assessment

The foregoing *asymptotic* analysis suggests somehow that, in general, the dependence of the truncation error on the dimension of the finite element space is very weak³. To corroborate this tentative generalization, in what follows, we numerically examine the sensitivity of the truncation error to mesh refinement. The employed reference grid, featuring $N_{g_0} = 1024$ points, together with the refined meshes, that contains $N_{g_1} = 3146$ and $N_{g_2} = 6712$ points ($\alpha_1 \approx 3$ and $\alpha_2 = 6.12$, respectively) are displayed in figure B.1. Using the training trajectories shown in figure 5.2, stress snapshots for each

¹Note that that $\|\bar{\mathbf{X}}\|_2 = \|\bar{\mathbf{S}}\|_2 \leq \|\mathbf{A}\|_2 \|\mathbf{X}\|_2 = \sqrt{\alpha N_g} \|\mathbf{S}\|_2$.

²In performing an uniform refinement with $\alpha \in \mathbb{N}$, the spatial patterns associated to the left singular vectors remains basically unaltered, allowing a fairer comparison between \mathbf{S} and $\bar{\mathbf{S}}$.

³This is in accordance with an observation made by Krysl et al. (2001): “While it may not be easy to come up with some general theoretical statements regarding M [number of sample points], numerical experiments indicate that the dependence of M on N [number of quadrature points of the finite element model] is very weak indeed”. It bears mentioning that Krysl et al. (2001) explicitly acknowledges the elusiveness of theoretically proving this property.

of these configurations were generated. Figure B.2 depicts the graph of the ratio between the computed singular values corresponding to the finer meshes with respect to their coarser counterparts, along with the asymptotic ratios, which, recall (see Eq.(B.2.3)), are given by:

$$\left. \frac{\lambda_i(\alpha N_g)}{\lambda_i(N_g)} \right|_{N_g \rightarrow \infty} = \sqrt{\alpha}. \quad (\text{B.3.1})$$




Figure B.1 Spatial discretizations used for studying the sensibility of the truncation error to mesh refinement. (a) $N_{g_0} = 1024$ (b) $N_{g_1} = 3146$ (c) $N_{g_2} = 6712$.

Observe that the two plots corresponding to the numerically calculated ratios exhibit similar patterns. They are, interestingly, coincident with their respective asymptotic values at the origin, following thereafter zigzag paths than run below these asymptotic lines for $i < 50$, approximately, and above for $i > 50$. It should be noted that deviations from the asymptotic values are not substantial, the maximum being barely 15%; this fact provides convincing proof that, as initially suggested, the *truncation error remains practically unchanged upon mesh refinement*. This can be better appreciated in the plot of the ratio between the truncation error estimate associated to the finer meshes with respect to their coarse mesh counterpart, displayed in figure B.3.

B.3.1 Discussion of results

B.3.1.1 Case $\bar{\lambda}_i/\lambda_i = \sqrt{\alpha}$

A feature of the graphs contained in figure B.2 that merits further attention is the coincidence of the numerical and asymptotic graphs at the origin ($i = 1$). To explain this intriguing feature, we shall follow the same reasoning used in the asymptotic case. The only difference is that now the ratio between singular values before and after mesh refinement is not constant and, consequently, the relation between the corresponding snapshot matrices \mathbf{X} and $\bar{\mathbf{X}}$ cannot be established via a single linear operator \mathbf{A} . This conceptual difficulty can be easily overcome by, first, decomposing the snapshot matrices as a weighted sum of ordered, *separable* matrices:

$$\mathbf{X} = \sum_i \mathbf{X}^{(i)} = \sum_i \lambda_i \mathbf{U}_i \mathbf{V}_i^T \quad (\text{B.3.2})$$

$$\bar{\mathbf{X}} = \sum_i \bar{\mathbf{X}}^{(i)} = \sum_i \bar{\lambda}_i \bar{\mathbf{U}}_i \bar{\mathbf{V}}_i^T. \quad (\text{B.3.3})$$

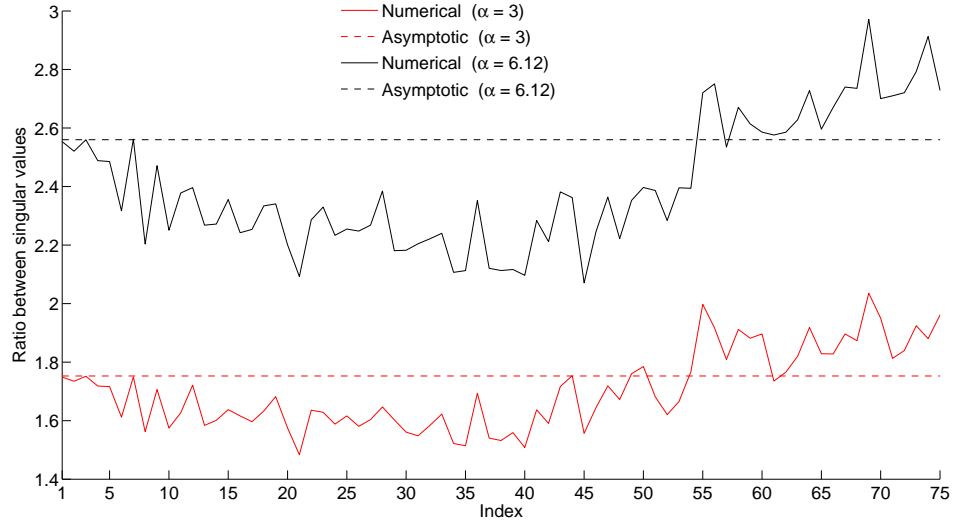


Figure B.2 Ratio between the singular values of the stress snapshot matrices obtained with the mesh layouts shown in B.1.b ($\alpha \approx 3$) and B.1.c ($\alpha = 6.12$) and the singular values corresponding to the coarser grid (figure B.1.a). The dashed lines indicate the (constant) ratio predicted by the asymptotic analysis ($\sqrt{\alpha}$)




Figure B.3 Ratio between the error bound estimate Eq.(B.1.2) associated to the meshes shown in B.1.b ($\alpha \approx 3$) and B.1.c ($\alpha = 6.12$) with respect to the one calculated with the coarser grid (figure B.1.a).

Next, in analogy to the asymptotic case, we introduce a certain matrix $\mathbf{A}^{(i)}$ connecting each pair of modal snapshots $\{\mathbf{X}^{(i)}, \bar{\mathbf{X}}^{(i)}\}$:

$$\bar{\mathbf{X}}^{(i)} = \mathbf{A}^{(i)} \mathbf{X}^{(i)} \quad (\text{B.3.4})$$

Taking the Frobenius norm of both sides of the above equation, one finally gets

$$\|\bar{\mathbf{X}}^{(i)}\|_2 = \bar{\lambda}_i = \|\mathbf{A}^{(i)} \mathbf{X}^{(i)}\|_2 \leq \|\mathbf{A}^{(i)}\|_2 \|\mathbf{X}^{(i)}\|_2 = \|\mathbf{A}^{(i)}\|_2 \lambda_i. \quad (\text{B.3.5})$$

The ratio between the singular values after and before refinement, thus, is

a function of the norm of the operator $\mathbf{A}^{(i)}$:

$$\frac{\bar{\lambda}_i}{\lambda_i} = f(\|\mathbf{A}^{(i)}\|_2) \leq \|\mathbf{A}^{(i)}\|_2. \quad (\text{B.3.6})$$

In particular, for $i = 1$, we have that

$$\frac{\bar{\lambda}_1}{\lambda_1} = f(\|\mathbf{A}^{(1)}\|_2) = \frac{\|\mathbf{A}^{(1)}\|_2}{\sqrt{N_g}} = \sqrt{\alpha}, \quad (\text{B.3.7})$$

On intuitive grounds, this can be taken as an indication that the operator $\mathbf{A}^{(i)}$ enjoys the same properties stipulated in the asymptotic case, i.e., the *spatial patterns associated to the first left singular vectors are essentially the same* for the three analyzed grids. An indirect manner of checking this is by comparing the *right* singular vectors \mathbf{V}_1 associated to each mesh; according to Eq.(A.4.5) the corresponding right singular vectors \mathbf{V}_1 should be identical in such cases. The plots of the right singular vectors, displayed in figure B.4, thus, clearly support our expectations, since the three graphs are undistinguishable.




Figure B.4 Plot of the components of the first right singular vector of the snapshot matrices associated to the three meshes shown in figure B.1.

B.3.1.2 Case $\bar{\lambda}_i/\lambda_i \neq \sqrt{\alpha}$

We can carry the above argument further to try to explain, albeit only vaguely, the overall trend exhibited by the numerical curves shown previously in figure B.2. For this purpose, we shall consider a more general scenario, in the sense that it embodies also the above discussed asymptotic case, in which the values of the variable corresponding to the i -th mode at each point of the refined mesh can be linearly interpolated from the respective values at neighboring points of the coarser mesh. Based on elementary interpolation properties — more specifically, the partition of unity property —, it can be shown that, in such cases:

$$\|\bar{\mathbf{X}}^{(i)}\|_2 \leq \sqrt{\alpha}\|\mathbf{X}^{(i)}\|_2 \Rightarrow \frac{\bar{\lambda}_i}{\lambda_i} \leq \sqrt{\alpha} \quad (\text{B.3.8})$$

(note that the equality holds in the limiting case $N_g \rightarrow 0$). Thus, this analysis indicates that the modes whose ratios happen to be below the asymptotic line (the first 50 modes, approximately) corresponds to this type of situations. Likewise, ratios greater than the asymptotic threshold implies that the above mentioned interpolatory proviso is not observed, and, therefore, as opposed to the case $\bar{\lambda}_i/\lambda_i \leq \sqrt{\alpha}$, the associated spatial modes before and after refinement are likely to exhibit significant differences.

Appendix C

Selection of sample points

This appendix is devoted to study more in depth an issue of paramount importance in the dimensionality reduction process, namely, the selection of representative *sample* points. To keep the exposition as simple as possible, the discussion is couched in terms of approximation of scalar-valued functions depending on scalar parameter — extension to multidimensional case is straightforward.

C.1 Introduction

Suppose that we wish to devise a method for the fast-run evaluation of a function $f = f(x, \mu)$, with $\mu \in \Omega_p = [\mu_0, \mu_F]$ (the parametric space) and $x \in \Omega = [x_0, x_F]$. To this end, we first perform a discretization of both Ω_p and Ω into T and N points, respectively. Then the function f is evaluated at each pair $\{x_i, \mu_j\}$, ($i = 1 \cdots N$, $j = 1 \cdots T$), and the resulting values stored in a $N \times T$ matrix, that will be denoted henceforth by \mathbf{F} ($F_{ij} = f(x_i, \mu_j)$) — this is the so-called *snapshot matrix*.

Next, the singular value decomposition is applied to \mathbf{F} :

$$\mathbf{F} = \mathbf{U}\mathbf{S}\mathbf{V}^T. \quad (\text{C.1.1})$$

We saw in appendix A.2 that the column vectors of \mathbf{U} — termed the *left singular vectors* — whose same-numbered *singular values* $\mathbf{S}_{jj} = \lambda_j$ are greater than zero constitute a *basis* for the range of \mathbf{F} . It was also explained that if the first M singular values happens to be significantly larger than the last $R - M$ ones (R denotes the rank of \mathbf{F}), then the range of \mathbf{F} can be accurately approximated by the span of the first M left singular vectors — the *dominant* modes. The matrix formed by these vectors (the *leading* or *dominant* modes) will be henceforth denoted by $\Phi \in \mathbb{R}^{N \times M}$, whereas $\Psi \in \mathbb{R}^{N \times R - M}$ will symbolize the remaining (or *trailing*) left singular vectors.

Suppose now that we know the values of the function f at $K \geq M$ points $\{\hat{x}_1, \hat{x}_2 \cdots \hat{x}_K\} \in \Omega$ (denoted collectively as $\hat{\mathbf{x}}$), and we are confronted with the task¹ of estimating, using the basis Φ provided by the SVD of

¹For the particular case of $K = M$ (number of points equal to the cardinality of the basis), this constitutes an *interpolation* problem; for $K > M$, strictly, the approximation can be no longer termed interpolation since the approximated function does not pass through the known values — *reconstruction* would be a more appropriate term.

\mathbf{F} , the values of the function at the remaining points in the domain. By approximating f as a linear combination of the basis vectors $\Phi_{\mathbf{i}}$ ($\mathbf{i} = \mathbf{1} \cdots \mathbf{M}$):

$$\mathbf{f}(\mathbf{x}, \mu) \approx \mathbf{f}^{\text{apr}}(\mathbf{x}, \mu) = \Phi(\mathbf{x})\mathbf{c}(\mu), \quad (\text{C.1.2})$$

this task boils down to evaluate the coefficient vector $\mathbf{c}(\mu) \in \mathbb{R}^M$ appearing in the above equation. For this purpose, we have at our disposal the following system of K equations with M unknowns:

$$\hat{\mathbf{f}} = \hat{\Phi}\mathbf{c} \quad (\text{C.1.3})$$

where $\hat{\mathbf{f}} = \mathbf{f}(\hat{\mathbf{x}}, \mu)$ stands for the vector containing the “known” values of f at the sample points, and $\hat{\Phi} = \Phi(\hat{\mathbf{x}})$ designates the sub-block matrix constructed by collecting the rows of Φ associated to such sample points. The system of equations represented by Eq.(C.1.3) is overdetermined; thus, only a solution in the “least-square sense” can be obtained: find $\mathbf{c} \in \mathbb{R}^M$ such that $\|\hat{\mathbf{f}} - \hat{\Phi}\mathbf{c}\|_2$ is minimum. This solution is given (see DeVore et al. (2001), p. 156) by the following vector of coefficients²:

$$\mathbf{c} = \overbrace{\left(\hat{\Phi}^T \hat{\Phi}\right)^{-1}}^{\mathbf{M}^{-1}} \hat{\Phi}^T \hat{\mathbf{f}}. \quad (\text{C.1.4})$$

It follows from the above equation that the coefficient vector \mathbf{c} can be only obtained if the inverse of $\mathbf{M} = \hat{\Phi}^T \hat{\Phi}$ exists; this condition, in turn, is satisfied whenever the columns of $\hat{\Phi}$ are linearly independent (full column rank).

C.2 Selection strategies

Consider now that we are given the freedom to choose the location of the abscissas $\{\hat{x}_1, \hat{x}_2 \cdots \hat{x}_K\} \in \Omega$ at which the values of the function are known. This situation is encountered, for instance, when one wishes to evaluate the integral of f over the domain Ω , for any value of μ , in an *online* cost-efficient manner:

$$\begin{aligned} I(\mu) &= \int_{x_0}^{x_F} f(x, \mu) dx \approx \int_{x_0}^{x_F} \Phi(\mathbf{x})\mathbf{c}(\mu) d\mathbf{x} \\ &= \sum_{k=0}^K \overbrace{\int_{x_0}^{x_F} \Phi(\mathbf{x}) d\mathbf{x}}^{\omega_k} \left(\hat{\Phi}^T \hat{\Phi}\right)^{-1} \hat{\Phi}_k^T \hat{\mathbf{f}}_k \\ &= \sum_{k=0}^K \omega_k f(\hat{x}_k). \end{aligned} \quad (\text{C.2.1})$$

The weighting coefficients ω_k does not depend on μ and can be, thereby, precomputed *offline*. Thus, as in classical Gaussian quadratures schemes, computation of the integral $I(\mu)$ only requires to evaluate the function at K distinct points. The question that arises naturally at this stage is which location of the sample points leads to minimum approximation errors. We discuss this crucial issue in in the sequel.

²The matrix $\left(\hat{\Phi}^T \hat{\Phi}\right)^{-1} \hat{\Phi}^T$ is known as the pseudo-inverse of $\hat{\Phi}$; notice that it coincides with $\hat{\Phi}^{-1}$ for $K = M$.

C.2.1 Statement of the optimization problem

In the limiting case of $K = N$, that is, when the entire set of points is used, Eq.(C.1.3) degenerates into the expression of the *orthogonal* projection of \mathbf{f} onto the space spanned by Φ :

$$\mathbf{f}^{best}(x, \mu) = \Phi \Phi^T \mathbf{f}. \quad (\text{C.2.2})$$

The superscript “best” in \mathbf{f}^{best} has been appended to indicate that this is the best approximation that can be achieved with the first M basis vectors provided by the SVD. To omit from the discussion discrepancies induced by deficient sampling of the parametric space, we shall examine in the following the error — with respect to the best approximation — committed in approximating each column of the snapshot matrix:

$$\mathbf{e}^{rec} = \mathbf{F}^{best} - \mathbf{F}^{apr}, \quad (\text{C.2.3})$$

where:

$$\mathbf{F}^{best} = \Phi \Phi^T \mathbf{F}, \quad (\text{C.2.4})$$

and

$$\mathbf{F}^{apr} = \Phi \left(\hat{\Phi}^T \hat{\Phi} \right)^{-1} \hat{\Phi}^T \hat{\mathbf{F}}. \quad (\text{C.2.5})$$

Here, $\hat{\mathbf{F}}$ designates the block matrix of \mathbf{F} formed by the rows corresponding to the selected points. By virtue of the single of value decomposition, the snapshot matrix can be decomposed into two mutually orthogonal components: one that lies in the span of Φ (the first M dominant left singular vectors), and another that resides in the span of Ψ (the trailing left singular vectors):

$$\mathbf{F} = \Phi \mathbf{S}_\Phi \mathbf{V}_\Phi^T + \Psi \mathbf{S}_\Psi \mathbf{V}_\Psi^T. \quad (\text{C.2.6})$$

Inserting Eq.(C.2.6), Eq.(C.2.5) and Eq.(C.2.4) into Eq.(C.2.3), and assuming that the inverse of \mathbf{M}^{-1} exists, we arrive, upon trivial manipulation, at:

$$\mathbf{e}^{rec} = \Phi \left(\hat{\Phi}^T \hat{\Phi} \right)^{-1} \hat{\Phi}^T \hat{\Psi} \mathbf{S}_\Psi \mathbf{V}_\Psi^T. \quad (\text{C.2.7})$$

The optimization problem associated to the determination of the sample points can be, thus, posed as follows: find $\hat{\mathbf{x}} \in \mathbb{R}^K$ as the minimizer of the (Frobenius) norm of the error defined³ in Eq.(C.2.7):

$$\hat{\mathbf{x}} = \arg \min_{\hat{\mathbf{z}}} \left\| \left(\hat{\Phi}^T(\hat{\mathbf{z}}) \hat{\Phi}(\hat{\mathbf{z}}) \right)^{-1} \hat{\Phi}^T(\hat{\mathbf{z}}) \hat{\Psi}(\hat{\mathbf{z}}) \mathbf{S}_\Psi \mathbf{V}_\Psi^T \right\|_2, \quad (\text{C.2.8})$$

subject to the constraints:

$$\hat{\mathbf{z}}_1 \in \Omega, \hat{\mathbf{z}}_2 \in \Omega, \dots, \hat{\mathbf{z}}_K \in \Omega. \quad (\text{C.2.9})$$

This is the *continuous* version of the optimization problem — each \hat{x}_k is allowed to vary within the interval ⁴ Ω . However, note that, in a general context, the values of the function may be only available at N points of

³Notice that, because of the orthonormality of Φ , the minimizer of $\|\Phi \mathbf{c}(\hat{\mathbf{z}})\|_2$ is equal to the minimizer of $\|\mathbf{c}(\hat{\mathbf{z}})\|_2$

⁴The “best-point” interpolation points method proposed by Nguyen et al. (2008) is based on the solution of this non-linear, least square minimization problem (they use a

the domain; thus, the problem posed in C.2.9 should be more appropriately viewed as a discrete (combinatorial) optimization problem: find the indices $\{p_1, p_2 \dots p_K\}$ as the minimizer of the following optimization problem:

$$\min_{\{r_1, r_2 \dots r_K\}} \left\| \left(\hat{\Phi}^T \hat{\Phi} \right)^{-1} \hat{\Phi}^T \hat{\Psi} \mathbf{S}_\Psi \mathbf{V}_\Psi^T \right\|_2, \quad (\text{C.2.10})$$

where $\hat{\Phi}^T$ and $\hat{\Psi}$ denotes the restricted or *gappy* matrices containing the $\{r_1, r_2 \dots r_K\}$ rows of the basis matrix Φ and the matrix of trailing left singular vectors Ψ , respectively.

C.2.2 Error bound

The matrix \mathbf{V}_Ψ^T appearing in Eq.(C.2.10) are formed by $R - M$ orthonormal column vectors; thus, we have that

$$\|\mathbf{V}_\Psi^T\|_2 = \sqrt{\mathbf{R} - \mathbf{M}}. \quad (\text{C.2.11})$$

Furthermore, \mathbf{S}_Ψ denotes a diagonal matrix containing the neglected singular values and, therefore, its Frobenius norm is given by:

$$\|\mathbf{S}_\Psi\|_2 = \sqrt{\sum_{i=\mathbf{M}+1}^{\mathbf{R}} \lambda_i^2}. \quad (\text{C.2.12})$$

Substitution of these expressions into the objective function in Eq.(C.2.10), and consideration of the general properties of norms, leads to the following upper bound for the approximation error:

$$\|\mathbf{e}^{\text{rec}}\|_2 \leq \sqrt{\mathbf{M}(\mathbf{R} - \mathbf{M}) \sum_{i=\mathbf{M}+1}^{\mathbf{R}} \lambda_i^2} \left\| \left(\hat{\Phi}^T \hat{\Phi} \right)^{-1} \hat{\Phi}^T \hat{\Psi} \right\|_2. \quad (\text{C.2.13})$$

Notice that factor $\sqrt{M(R - M) \sum_{i=M+1}^R \lambda_i^2}$ depends exclusively on the level of truncation of the basis and not of the particular choice of sample points⁵. The only term affected by the sampling of the points is the right-most factor in the above inequality, henceforth denoted by k^{obj} :

$$k^{\text{obj}} := \left\| \left(\hat{\Phi}^T \hat{\Phi} \right)^{-1} \hat{\Phi}^T \hat{\Psi} \right\|_2 \quad (\text{C.2.14})$$

Levenberg-Marquardt algorithm for this purpose). This strategy, however, is computationally costly and convoluted, since it requires the computation of the Jacobian matrices of the basis vectors Φ and Ψ , which might not be defined at certain points in the domain. Furthermore, in multidimensional problems, such as the one addressed in chapter 4, entries $F_{i,j}$ and $F_{i+1,j}$ may not correspond to spatially contiguous points in the domain, and hence, the objective function cannot be expected to vary continuously.

⁵This factor is intimately connected with the truncation error $\mathbf{F} - \mathbf{F}^{\text{best}}$. Indeed, according to the decomposition introduced in Eq.(C.2.6), we have that $\|\mathbf{F} - \mathbf{F}^{\text{best}}\|_2 =$

$$\|\Psi \mathbf{S}_\Psi \mathbf{V}_\Psi^T\|_2 \leq (\mathbf{R} - \mathbf{M}) \sqrt{\sum_{i=\mathbf{M}+1}^{\mathbf{R}} \lambda_i^2}.$$

It is interesting to note that this term only entails inner products “in the gappy sense” between both bases $\hat{\Phi}$ and $\hat{\Psi}$. Since $\|\hat{\Phi}^T \hat{\Psi}\|_2 \leq \sqrt{R - M}$, we can also assert that:

$$\|\mathbf{F}^{\text{apr}} - \mathbf{F}^{\text{best}}\|_2 \leq (\mathbf{R} - \mathbf{M}) \sqrt{\mathbf{M} \sum_{i=\mathbf{M}+1}^{\mathbf{R}} \lambda_i^2} \left\| \left(\hat{\Phi}^T \hat{\Phi} \right)^{-1} \right\|_2. \quad (\text{C.2.15})$$

Thus, the only term that can grow unboundedly depending on the choice is $\left\| \left(\hat{\Phi}^T \hat{\Phi} \right)^{-1} \right\|_2$.

C.2.3 Greedy algorithms

Barring the notable exception of Nguyen et al. (2008) and their “best points” interpolation method (see footnote ⁴), proposals found in the literature dealing with the choice of optimal sample points relies on the minimization of the upper bound error estimate shown in Eq.(C.2.15); this minimization problem can be formally stated as follows: find indices $\{p_1, p_2 \cdots p_K\}$ such that

$$\{p_1, p_2 \cdots p_K\} = \arg \min_{\{r_1, r_2 \cdots r_K\}} \left\| \left(\hat{\Phi}^T \hat{\Phi} \right)^{-1} \right\|_2. \quad (\text{C.2.16})$$

Although considerably simpler than the problem stated in (C.2.9), the solution of Eq.(C.2.10) stills poses a major hurdle due essentially to its *combinatorial* character. An exhaustive search is in general not viable because of the sheer number of feasible solutions ($\binom{N}{K}$); for this reason, one has to resort to *heuristic* methods able to determine, at least, *sub-optimal* solutions.

The approaches proposed in the literature are all based on the so-called *greedy paradigm*. A greedy method is any algorithm that solves the problem by making the *locally optimal choice* at each stage with the hope of finding the global optimum. Accordingly, applied to the case under consideration, a greedy algorithm would select the first index as

$$p_1 = \arg \min_{1 \leq r_1 \leq N} \left\| \left(\Phi_{1,r_1} \Phi_{1,r_1} \right)^{-1} \right\|_2 = \arg \max_{r_1} |\Phi_{1,r_1}|, \quad (\text{C.2.17})$$

that is, the index corresponding to the point at which the first basis vector $\Phi_{\mathbf{1}}$ reaches its maximum (in absolute value). Likewise, the k – *th* index⁶ would be determined⁷ by the condition:

$$p_k = \arg \min_{1 \leq r_k \leq N} \left\| \left(\hat{\mathbf{D}}_{k-1}^T \hat{\mathbf{D}}_{k-1} \right)^{-1} \right\|_2. \quad (\text{C.2.18})$$

⁶This greedy strategy is only applicable for the first M indices. When $K > M$, there are no available bases to match the indices $p_{M+1}, p_{M+2} \cdots p_K$. One plausible, ad hoc route to circumvent this deficiency may be to “re-use” as basis vectors for $p_{M+1}, p_{M+2} \cdots$ the sequence $\Phi_{\mathbf{1}}, \Phi_{\mathbf{2}} \cdots$.

⁷The qualifier “greedy” refers to the fact that the k – *th* index is determined taking into account only the “current” basis vector $\Phi_{\mathbf{k}}$ and the previously fixed indices $\{p_1, p_2 \cdots p_{k-1}\}$; the algorithm, thus, proceeds *greedily*, i.e., it only considers “local short-term gains” in each selection, with no regard for the future consequences of such choices. The reader wishing for further coverage of greedy algorithms is referred to Lovasz et al. (2003) and Parhami (1999). The former reference provides a compelling illustration of the performance (often misperformance) of greedy algorithms in a famous task in discrete optimization: the Traveling Salesman Problem.

where $\hat{\mathbf{D}}_{k-1}$ denotes the restricted or “gappy” basis matrix formed by rows $\{p_1, p_2 \dots p_{k-1}\}$ of the matrix containing the first $k-1$ basis vectors ($\mathbf{D}_{k-1} = [\Phi_1, \Phi_2 \dots \Phi_{k-1}]$).

Based on the observation that all norms are equivalent, Willcox (2006) and Astrid et al. (2008) advocate to adopt as objective function the condition number of matrix $\hat{\Phi}^T \hat{\Phi}$ rather than the Frobenius norm of its inverse. On the other hand, in the algorithm proposed by Chaturantabut and Sorensen (2010) — called the Discrete Empirical Interpolation method (DEIM)⁸ —, the k -th point is chosen as that at which the difference, in absolute value, between the basis vector Φ_k and its approximation from interpolation at points $x_{p_1}, x_{p_2} \dots x_{p_{k-1}}$ using basis $\mathbf{D}_{k-1} = [\Phi_1, \Phi_2 \dots \Phi_{k-1}]$ is maximum:

$$p_k = \arg \max_{1 \leq r_k \leq N} |\Phi_k - (\hat{\mathbf{D}}_{k-1}^T \hat{\mathbf{D}}_{k-1})^{-1} \hat{\mathbf{D}}_{k-1}^T \hat{\Phi}_k|, \quad (\text{C.2.19})$$

where⁹ the matrix $\hat{\Phi}_k$ stands for the vector formed by the p_1 -th, p_2 -th \dots p_{k-1} -th rows of Φ_k .

C.3 Example

This section is intended to illustrate the performance of the various greedy strategies described in the foregoing. For this purpose, consider the following function:




Figure C.1 Plot of the parametric function defined in Eq.(C.3.1) for several values of the parameter μ .

$$f(x, \mu) = \begin{cases} 3 \cdot 10^{-2} \mu^3 x^2 \sin(x\mu) & x \in [0, 4[\\ 3 \cdot 10^{-4} \mu^3 x^4 & x \in [4, 9[\\ \sin(3x\mu) + \cos(2x\mu) & x \in [9, 9.5[\\ \sin(3x\mu) + \cos(x\mu) & x \in [9.5, 10] \end{cases} \quad (\text{C.3.1})$$

⁸The DEIM is the discrete version of the Empirical Interpolation Method, pioneered by Barrault et al. (2004) (see section 1.2.2.1).

⁹Note that for $k \leq M$, $\hat{\mathbf{D}}_{k-1}$ is a square matrix and, thereby, the term $(\hat{\mathbf{D}}_{k-1}^T \hat{\mathbf{D}}_{k-1})^{-1} \hat{\mathbf{D}}_{k-1}^T$ reduces to $\hat{\mathbf{D}}_{k-1}^{-1}$.

where¹⁰ $\mu \in \Omega_p = [-1, 1]$. Both Ω_p and Ω are discretized into $T = 100$ and $N = 500$ equally spaced points, respectively. The snapshot matrix resulting from evaluating the function at each pair $\{x_i, \mu_j\}$ ($[\mathbf{F}]_{ij} = \mathbf{f}(\mathbf{x}_i, \mu_j)$) is then subjected to the Single Value Decomposition $\mathbf{F} = \mathbf{U}\mathbf{S}\mathbf{V}^T$. To discriminate which left singular vectors are “dominant” (in a statistical sense), we examine the plot of the singular values $S_{ii} = \lambda_i$ ($i = 1, 2, \dots$), displayed in figure C.2. The magnitude of the singular values becomes vanishingly small at around $i = 10$; this tells us than taking, say, the first $M = 10$ ($\Phi = [\mathbf{U}_1, \mathbf{U}_2 \dots \mathbf{U}_{10}]$) modes suffices to replicate with reasonable accuracy the original snapshots.




Figure C.2 Magnitude of the singular values.

Another (more sound) way of determining an appropriate level of truncation is to analyze the graph of the following bound for the *truncation* error $\mathbf{F} - \mathbf{F}^{\text{best}}$:

$$\|\mathbf{F} - \mathbf{F}^{\text{best}}\|_2 = \|\Psi\mathbf{S}\Psi^T\mathbf{V}\|_2 \leq (R - M) \sqrt{\sum_{i=M+1}^R \lambda_i^2}. \quad (\text{C.3.2})$$

where $R = 25$ stands for the rank of \mathbf{F} . This expression follows easily from equations (C.2.6) and (C.2.12). Figure C.3 contains the representation of the above *a priori* error estimator (divided by $\|\mathbf{F}\|_2$):

$$e_b^{\text{trun}}(M) = (R - M) \frac{\sqrt{\sum_{i=M+1}^R \lambda_i^2}}{\|\mathbf{F}\|_2}. \quad (\text{C.3.3})$$

Inspection of this plot indicates that, in taking $M = 10$, the approximation error will be below 1%.

¹⁰In order to assess the accuracy of the approximation in extreme scenarios, the function has been purposefully crafted to exhibit discontinuities at $x = 4$, $x = 9$ and $x = 9.5$ (see figure C.1).

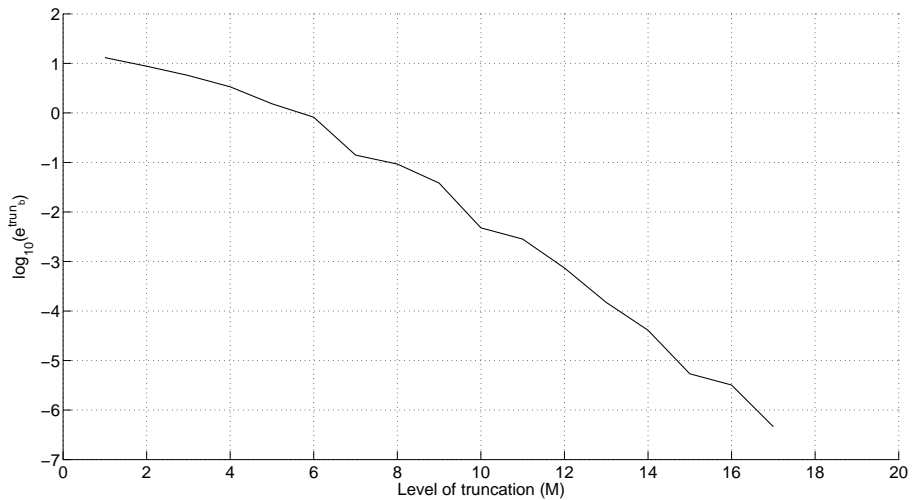


Figure C.3 Logarithm of the error bound defined in Eq.(C.3.2) versus the level of truncation M .

C.3.1 Comparison between greedy algorithms

Figure C.4 contains the plots of the error estimate $\frac{\|\mathbf{F} - \mathbf{F}^{\text{appr}}(\mathbf{M}, \mathbf{K})\|_2}{\|\mathbf{F}\|_2}$ against the number of modes M for the case in which $K = M$ (interpolation), and for four distinct selection procedures. The curves labeled as “DEIM”, “WILL-COX” and “INV” correspond to *greedy* algorithms based on the proposal by Chaturantabut and Sorensen (2010) (Eq.(C.2.19)), on the minimization of the condition number (Willcox (2006)), and on the minimization of $\|\mathbf{M}^{-1}\|_2$ (Eq.(C.2.12)), respectively. The graph with the “BRUTE” label, on the other hand, is not connected to any greedy algorithm, but rather represents the result of a “brute force” search: for each M , 10^6 random combinations of indices were generated and their associated error estimates evaluated; the “BRUTE” graph is the plot, for each M , of the minima of such estimates. Finally, for completeness, the plot of the truncation error measure $\|\mathbf{F} - \mathbf{F}^{\text{best}}\|_2$ is also displayed (labeled as “BEST”, since it constitutes a lower bound for the approximation error.)

Three things deserve notice in figure C.4:

- Differences in performance between the greedy algorithms DEIM, WILL-COX and INV are not substantial. For moderate levels of truncation, the DEIM seems to yield more accurate results than the other two, but as M increases, differences becomes imperceptible. This observation would make the selection procedure proposed by Chaturantabut and Sorensen (2010) (DEIM) the most attractive strategy, since it is, by far, the most economic method in terms of computational operations.
- The “brute force” approach yields the most accurate approximation for $M \in [4, 6]$; the error estimate for $M = 5$ is 7.9%, only 0.2% above the truncation error, while discrepancies with greedy methods ranges between 12% and 13%. This provides convincing proof that the sets of indices constructed via greedy algorithms are only sub-optimal. Obviously, the brute force approach becomes impractical as N and M increase.

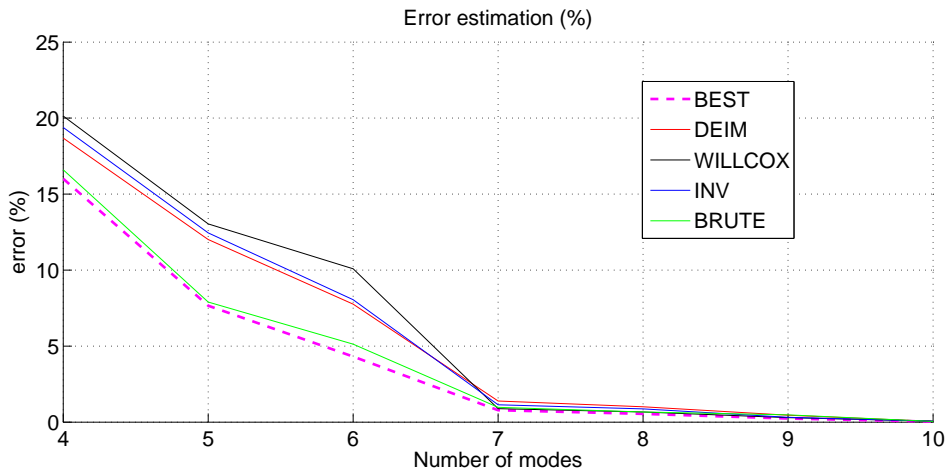


Figure C.4 Error measure $\frac{\|\mathbf{F} - \mathbf{F}^{\text{appr}}(M)\|_2}{\|\mathbf{F}\|_2} 100$ versus the level of truncation M , for the case in which $K = M$ (interpolation) and for four distinct selection procedures (labeled as DEIM, WILLCOX, INV and BRUTE). The graph denoted as “BEST” corresponds to the truncation error estimate $\frac{\|\mathbf{F} - \mathbf{F}^{\text{best}}(M)\|_2}{\|\mathbf{F}\|_2} 100$.

- The approximation error curves DEIM, WILLCOX, INV and BRUTE mimic essentially the pattern of the truncation error graph (BEST). This can be also inferred from expression (C.2.15). Indeed, it takes elementary algebra to show from this expression that:

$$\|\mathbf{F} - \mathbf{F}^{\text{appr}}\|_2 = \sqrt{1 + \alpha \frac{M}{R - M} k^{\text{obj}^2}} \|\mathbf{F} - \mathbf{F}^{\text{best}}\|_2 \quad (\text{C.3.4})$$

where k^{obj} was defined in Eq.(C.2.14), and $\alpha > 0$. For completeness, the value of the factor C in the above expression as a function of M , and for each sampling strategy, is displayed in figure C.5. Note that this factor barely varies with increasing M — it oscillates between 1.5 and 3.5 for all four sampling methods.

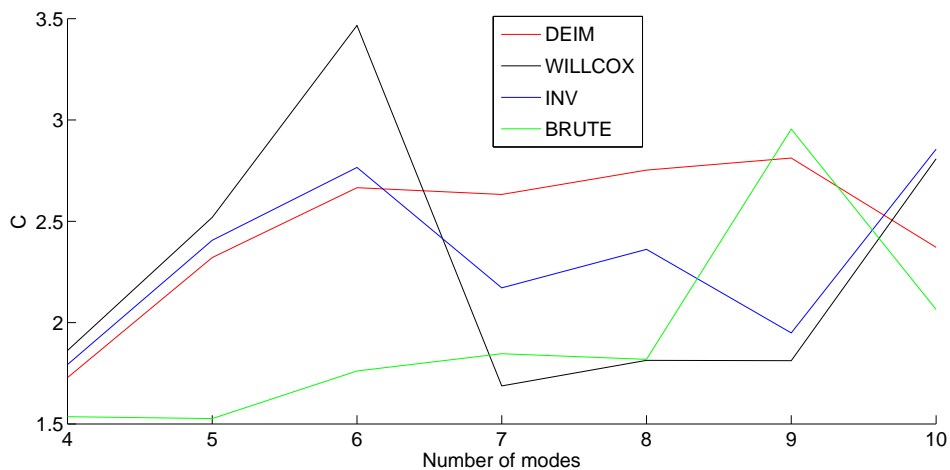


Figure C.5 Factor C defined in Eq.(C.3.4) versus the level of truncation M .

Observation C.3.1 *A far-reaching conclusion can be drawn from the above consideration. The fact that variable C in Eq.(C.3.4) does not exhibit a clear decreasing tendency, but rather oscillates between 1.5 and 3.5 for increasing values of M suggests that, if the approximation error $\|\mathbf{F} - \mathbf{F}^{\text{apr}}\|_2$ is to be significantly reduced, it proves more effective to reduce the truncation error by increasing the number of basis vectors¹¹ rather than by struggling in elaborating sophisticated sampling procedures, which only would affect the magnitude of such factor.*

The locations of the first six points for the algorithms labeled previously as INV, WILLCOX and DEIM, together with the corresponding basis functions, are depicted in figure C.6. It is instructive to note that the DEIM greedy algorithm invariably places the k -th point at, or close to, the global maximum of the k -th basis function¹². The points provided by the “INV” algorithm — the one based in the greedy minimization of $\|\mathbf{M}^{-1}\|_2$ — are also located in the vicinity of the global maxima (in fact, they are coincident with DEIM points for $k = 1, 2, 4, 6$). The placement of the sampling points furnished by the WILLCOX method, on the other hand, coincides with the other methods for¹³ $k = 1, 4$, but are notably different in the other steps.

The placement of the DEIM points in relation with the critical points of the basis functions can be better appreciated in the overall representation of figure C.7. For comparison purposes, we have also plotted in this figure the points calculated by the “brute” force method — which, according to the graph shown previously in figure C.4, provides, for this level of truncation ($M = 6$), a better approximation than the DEIM. Observe that the locations calculated with the “brute force” strategy do not lie necessarily around the the global maxima or other critical points of the basis functions; this is because, as opposed to the DEIM points, such points depends not only on the *dominant* modes $\Phi = [\Phi_1 \Phi_2 \cdots \Phi_6]$, but also on the topology of the *trailing* basis functions Ψ .

¹¹Recall that the error $\|\mathbf{F} - \mathbf{F}^{\text{best}}\|_2$ patterns somehow the decreasing tendency of the singular values graph.

¹²It may be interesting in future developments to further investigate this peculiarity, and to ascertain if a topological meaning can be definitely ascribed to the points determined by the Discrete Empirical Interpolation Method.

¹³Since for $k = 1$, matrix \mathbf{M} is in fact a scalar, all points in the first greedy step are equally optimal in terms of condition number. This is why the WILLCOX method is usually accompanied by a *screening* criterion (Astrid et al., 2008) that serves to select an optimal point in such first step. The one implemented here is to take as first best point that at which $|\Phi_1|$ is maximum. Incidentally, this explains why the first points furnished by the three methods (INV, WILLCOX and DEIM) are coincident.

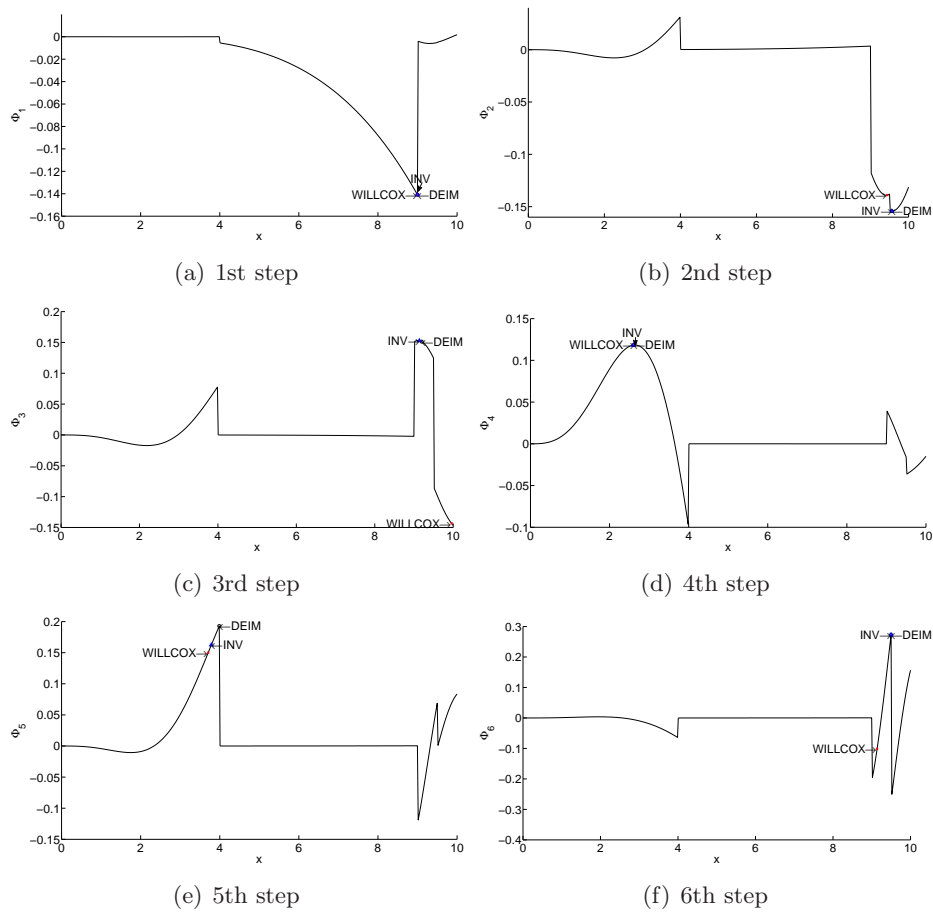


Figure C.6 Six first SVD basis functions together with the corresponding sampling points selected by the “INV”, “WILLCOX” and “DEIM” algorithms at each step of the greedy process.




Figure C.7 Six first basis functions together with the points calculated with the DEIM (circles) and the brute force algorithm (crosses).

References

- Amsallem, D., Cortial, J., Carlberg, K., Farhat, C., 2009. A method for interpolating on manifolds structural dynamics reduced-order models. *International Journal for Numerical Methods in Engineering* 80 (9), 1241–1258.
- An, S., Kim, T., James, D., 2009. Optimizing cubature for efficient integration of subspace deformations. *ACM transactions on graphics* 27 (5), 165.
- Astrid, P., 2004. Reduction of process simulation models: a proper orthogonal decomposition approach. Technische Universiteit Eindhoven.
- Astrid, P., Weiland, S., Willcox, K., Backx, T., 2008. Missing point estimation in models described by proper orthogonal decomposition. *Automatic Control, IEEE Transactions on* 53 (10), 2237–2251.
- Barrault, M., Maday, Y., Nguyen, N., Patera, A., 2004. An empirical interpolation method: application to efficient reduced-basis discretization of partial differential equations. *Comptes Rendus Mathematique* 339 (9), 667–672.
- Belytschko, T., Liu, W. K., Moran, B., 2001. *Nonlinear Finite Elements for Continua and Structures*. John Wiley and Sons Ltd., New York.
- Boyaval, S., 2007. Reduced-basis approach for homogenization beyond the periodic setting. Arxiv preprint math/0702674.
- Bui-Thanh, T., 2007. Model-constrained optimization methods for reduction of parameterized large-scale systems. Ph.D. thesis, Citeseer.
- Bui-Thanh, T., Willcox, K., Ghattas, O., 2008. Model reduction for large-scale systems with high-dimensional parametric input space. *SIAM Journal on Scientific Computing* 30 (6), 3270–3288.
- Bui-Thanh, T., Willcox, K., Ghattas, O., van Bloemen Waanders, B., 2007. Goal-oriented, model-constrained optimization for reduction of large-scale systems. *Journal of Computational Physics* 224 (2), 880–896.
- Carlberg, K., Bou-Mosleh, C., Farhat, C., 2011. Efficient non-linear model reduction via a least-squares petrov–galerkin projection and compressive tensor approximations. *International Journal for Numerical Methods in Engineering* 86 (2), 155–181.

- Carlberg, K., Farhat, C., 2008. A Compact Proper Orthogonal Decomposition Basis for Optimization-Oriented Reduced-Order Models. AIAA Paper 5964, 10–12.
- Chatterjee, A., 2000. An introduction to the proper orthogonal decomposition. *Current science* 78 (7), 808–817.
- Chaturantabut, S., Sorensen, D., 2010. Discrete empirical interpolation for nonlinear model reduction. In: *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on.* IEEE, pp. 4316–4321.
- Chen, W., Fish, J., 2006. A generalized space-time mathematical homogenization theory for bridging atomistic and continuum scales. *International Journal for Numerical Methods in Engineering* 67 (2), 253–271.
- Cortelezzi, L., Speyer, J., 1998. Robust reduced-order controller of laminar boundary layer transitions. *Physical Review E* 58 (2), 1906–1910.
- de Souza Neto, E., Feijóo, R., 2006. Variational foundations of multi-scale constitutive models of solid: small and large strain kinematical formulation. LNCC Research & Development Report 16.
- DeVore, R., Iserles, A., Suli, E., 2001. *Foundations of computational mathematics.* Cambridge Univ Pr.
- Dolbow, J., Khaleel, M., Mitchell, J., (US), P. N. N. L., of Energy, U. S. D., 2004. Multiscale mathematics initiative: a roadmap. Pacific Northwest National Laboratory.
- Dvorak, G., Wafa, A., Bahei-El-Din, Y., 1994. Implementation of the transformation field analysis for inelastic composite materials. *Computational Mechanics* 14 (3), 201–228.
- Efendiev, Y., Hou, T., 2009. *Multiscale finite element methods: theory and applications.* Springer Verlag.
- Everson, R., Sirovich, L., 1995. Karhunen–Loeve procedure for gappy data. *Journal of the Optical Society of America A* 12 (8), 1657–1664.
- Fish, J., 2009. *Multiscale Methods: Bridging the Scales in Science and Engineering.* Oxford Univ Pr.
- Fish, J., Chen, W., 2004. Space-time multiscale model for wave propagation in heterogeneous media. *Computer Methods in applied mechanics and engineering* 193 (45-47), 4837–4856.
- Fish, J., Chen, W., Li, R., 2007. Generalized mathematical homogenization of atomistic media at finite temperatures in three dimensions. *Computer methods in applied mechanics and engineering* 196 (4-6), 908–922.
- Fish, J., Chen, W., Nagai, G., 2002. Non-local dispersive model for wave propagation in heterogeneous media: one-dimensional case. *International Journal for Numerical Methods in Engineering* 54 (3), 331–346.

- Fish, J., Shek, K., 1999. Finite deformation plasticity for composite structures: computational models and adaptive strategies. *Computer Methods in Applied Mechanics and Engineering* 172 (1-4), 145–174.
- Fish, J., Shek, K., Pandheeradi, M., Shephard, M., 1997. Computational plasticity for composite structures based on mathematical homogenization: Theory and practice. *Computer Methods in Applied Mechanics and Engineering* 148 (1-2), 53–73.
- Galbally, D., Fidkowski, K., Willcox, K., Ghattas, O., 2010. Non-linear model reduction for uncertainty quantification in large-scale inverse problems. *International Journal for Numerical Methods in Engineering* 81 (12), 1581–1608.
- Ganapathysubramanian, S., Zabaras, N., 2004. Design across length scales: a reduced-order model of polycrystal plasticity for the control of microstructure-sensitive material properties. *Computer Methods in Applied Mechanics and Engineering* 193 (45-47), 5017–5034.
- Germain, P., 1973. The method of virtual power in continuum mechanics. Part 2: Microstructure. *SIAM Journal on Applied Mathematics* 25 (3), 556–575.
- Giusti, S., Blanco, P., de Souza Neto, E., Feijóo, R., 2009. An assessment of the Gurson yield criterion by a computational multi-scale approach. *Engineering Computations* 26 (3), 281–301.
- Grepl, M., Maday, Y., Nguyen, N., Patera, A., 2007. Efficient reduced-basis treatment of nonaffine and nonlinear partial differential equations. *Mathematical Modelling and Numerical Analysis* 41 (3), 575–605.
- Gross, D., Seelig, T., 2011. *Fracture mechanics: with an introduction to micromechanics*. Springer.
- Gurson, A., et al., 1977. Continuum theory of ductile rupture by void nucleation and growth: Part I-Yield criteria and flow rules for porous ductile media. *Journal of Engineering Materials and Technology* 99 (1), 2–15.
- Hill, R., 1963. Elastic properties of reinforced solids: some theoretical principles. *Journal of the Mechanics and Physics of Solids* 11 (5), 357–372.
- Huynh, D., Patera, A., 2007. Reduced basis approximation and a posteriori error estimation for stress intensity factors. *International Journal for Numerical Methods in Engineering* 72 (10), 1219–1259.
- Kerschen, G., Golinval, J., Vakakis, A., Bergman, L., 2005. The method of proper orthogonal decomposition for dynamical characterization and order reduction of mechanical systems: an overview. *Nonlinear Dynamics* 41 (1), 147–169.
- Krysl, P., Lall, S., Marsden, J., 2001. Dimensional model reduction in nonlinear finite element dynamics of solids and structures. *International Journal for Numerical Methods in Engineering* 51 (4), 479–504.

- Landauer, T., Laham, D., Foltz, P., 1998. Learning human-like knowledge by singular value decomposition: A progress report. *Advances in neural information processing systems*, 45–51.
- Lee, J., Verleysen, M., 2007. *Nonlinear dimensionality reduction*. Springer Verlag.
- Lopez, R., 2008. *Neural networks for variational problems in engineering*. Ph.D. thesis, Technical University of Cataluña, Barcelona.
- Lopez, R., Balsa-Canto, E., Oñate, E., 2008. Neural networks for variational problems in engineering. *International Journal for Numerical Methods in Engineering* 75, 1341–1360.
- Lovasz, L., Pelikan, J., Vesztegombi, K., 2003. *Discrete Mathematics: Elementary and Beyond*. Springer.
- Michel, J., Moulinec, H., Suquet, P., 1999. Effective properties of composite materials with periodic microstructure: a computational approach. *Computer methods in applied mechanics and engineering* 172 (1-4), 109–143.
- Miehe, C., Schotte, J., Lambrecht, M., 2002. Homogenization of inelastic solid materials at finite strains based on incremental minimization principles. Application to the texture analysis of polycrystals. *Journal of the Mechanics and Physics of Solids* 50 (10), 2123–2167.
- Miehe, C., Schotte, J., Schroder, J., 1999. Computational micro-macro transitions and overall moduli in the analysis of polycrystals at large strains. *Computational Materials Science* 16 (1-4), 372–382.
- Monteiro, E., Yvonnet, J., He, Q., 2008. Computational homogenization for nonlinear conduction in heterogeneous materials using model reduction. *Computational Materials Science* 42 (4), 704–712.
- Nemat-Nasser, S., 1999. Averaging theorems in finite deformation plasticity. *Mechanics of Materials* 31 (8), 493–523.
- Nguyen, N., 2007. A posteriori error estimation and basis adaptivity for reduced-basis approximation of nonaffine-parametrized linear elliptic partial differential equations. *Journal of Computational Physics* 227 (2), 983–1006.
- Nguyen, N., 2008. A multiscale reduced-basis method for parametrized elliptic partial differential equations with multiple scales. *Journal of Computational Physics* 227 (23), 9807–9822.
- Nguyen, N., Patera, A., Peraire, J., 2008. A best points interpolation method for efficient approximation of parametrized functions. *Int. J. Numer. Meth. Engng* 73, 521–543.
- Nguyen, N., Peraire, J., 2008. An efficient reduced-order modeling approach for non-linear parametrized partial differential equations. *International Journal for Numerical Methods in Engineering* 76 (1), 27–55.

- Oskay, C., Fish, J., 2007. Eigendeformation-based reduced order homogenization for failure analysis of heterogeneous materials. *Computer Methods in Applied Mechanics and Engineering* 196 (7), 1216–1243.
- Parhami, B., 1999. *Introduction to parallel processing: algorithms and architectures*. Vol. 1. Springer.
- Press, W., Flannery, B., Teukolsky, S., Vetterling, W., 1993. *Numerical Recipes in Fortran 77: The Art of Scientific Computing* 1. Cambridge: Cambridge University Press.
- Quarteroni, A., Sacco, R., Saleri, F., 2000. *Numerical Mathematics*. Springer, New York.
- Reddy, B. D., 1998. *Introductory Functional Analysis*. Springer-Vedag, New York.
- Rowley, C., 2005. Model reduction for fluids, using balanced proper orthogonal decomposition. To appear in *Int. J. on Bifurcation and Chaos*.
- Rowley, C., Colonius, T., Murray, R., 2004. Model reduction for compressible flows using pod and galerkin projection. *Physica D: Nonlinear Phenomena* 189 (1-2), 115–129.
- Rozza, G., Huynh, D., Patera, A., 2007. Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations. *Archives of Computational Methods in Engineering* 15 (3), 1–47.
- Ryckelynck, D., 2005. A priori hyperreduction method: an adaptive approach. *Journal of computational physics* 202 (1), 346–366.
- Sanchez-Palencia, E., 1974. Comportements local et macroscopique d'un type de milieux physiques heterogenes. *International Journal of Engineering Science* 12 (4), 331–351.
- Sanchez-Palencia, E., 1980. Non-homogeneous media and vibration theory. In: *Non-Homogeneous Media and Vibration Theory*. Vol. 127.
- Smith, L., 2002. A tutorial on principal components analysis. Cornell University, USA 51, 52.
- Šolín, P., 2006. *Partial differential equations and the finite element method*. Wiley-Blackwell.
- Stewart, G., 1993. On the early history of the singular value decomposition. *Siam Review* 35 (4), 551–566.
- Terada, K., Saiki, I., Matsui, K., Yamakawa, Y., 2003. Two-scale kinematics and linearization for simultaneous two-scale analysis of periodic heterogeneous solids at finite strain. *Computer Methods in Applied Mechanics and Engineering* 192 (31-32), 3531–3563.
- Tu, X., Andrade, J., Chen, Q., 2009. Return mapping for nonsmooth and multiscale elastoplasticity. *Computer Methods in Applied Mechanics and Engineering* 198 (30-32), 2286–2296.

- Vaseghi, S., 2007. Multimedia signal processing: theory and applications in speech, music and communications. Wiley.
- Vernerey, F., Liu, W., Moran, B., 2007. Multi-scale micromorphic theory for hierarchical materials. *Journal of the Mechanics and Physics of Solids* 55 (12), 2603–2651.
- Willcox, K., 2006. Unsteady flow sensing and estimation via the gappy proper orthogonal decomposition. *Computers & fluids* 35 (2), 208–226.
- Xiao, S., Belytschko, T., 2004. A bridging domain method for coupling continua with molecular dynamics. *Computer methods in applied mechanics and engineering* 193 (17-20), 1645–1669.
- Yuan, Z., Fish, J., 2008. Towards Realization of Computational Homogenization in Practice1. *Int. J. Numer. Meth. Engng* 73, 361–380.
- Yuan, Z., Fish, J., 2009. Multiple scale eigendeformation-based reduced order homogenization. *Computer Methods in Applied Mechanics and Engineering* 198 (21-26), 2016–2038.
- Yvonnet, J., He, Q., 2007. The reduced model multiscale method (R3M) for the non-linear homogenization of hyperelastic media at finite strains. *Journal of Computational Physics* 223 (1), 341–368.