Journal of Computational Physics 490 (2023) 112319

Contents lists available at ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

MWCAWE: A multivariate WCAWE approach for parametric model order reduction, and a sampling strategy for the bivariate case

Romain Rumpler [a,b,*], Quirin Aumann [c]

[a] The Marcus Wallenberg Laboratory for Sound and Vibration Research (MWL), Department of Engineering Mechanics, KTH Royal Institute of Technology, Stockholm, SE-100 44, Sweden
[b] Centre for ECO2 Vehicle Design, KTH Royal Institute of Technology, Stockholm, SE-100 44, Sweden
[c] Max Planck Institute for Dynamics of Complex Technical Systems, Sandtorstraße 1, 39106 Magdeburg, Germany

ARTICLE INFO

Article history:
Received 24 August 2022
Received in revised form 12 June 2023
Accepted 20 June 2023
Available online 3 July 2023

Dataset link: https://doi.org/10.5281/zenodo.7806380

Keywords:
Parametric model order reduction
Moment-matching method
Finite element method
Padé approximant
Vibro-acoustics

ABSTRACT

In this paper, the multivariate well-conditioned asymptotic waveform evaluation (MWCAWE) algorithm is presented. It is based on a univariate version of the algorithm, previously introduced for fast frequency sweeps. Taking advantage of the robustness of this univariate algorithm, the MWCAWE enables to effectively generate multi-parameter reduced-order models. Additionally, a residue-based contour following approach is introduced for a two-stage generation of a projection basis, here applied on bi-variate acoustic problems where frequency and a material parameter are in focus. The effectiveness of the proposed method is demonstrated on two poro-acoustic problems, highlighting both the good convergence property of the algorithm and the potential of the multi-patch strategy for parametric applications.

1. Introduction

Many problems in optimisation, design, and control require the evaluation of numerical models depending on multiple parameters during the early design stage and operation. Such parameters may include the driving frequency, material properties, or geometric features, among others. For example, dynamical systems are often evaluated using frequency sweep methods requiring the repeated evaluation of a typically large system of equations for a changing driving frequency. To allow an efficient simulation process, surrogate models, approximating the input-output behaviour of the original system while retaining the dependency on certain model parameters, are required. For this purpose, component-wise expansion techniques allowing to target selected degrees of freedom of interest, for example based on rational polynomial expansion in the form of Padé approximants, may be implemented; see [1,2] for univariate examples and [3,4] for a multivariate extension. Alternatively, projection-based parametric model order reduction (pMOR) has been intensely investigated and is now a well-established approach to obtain such surrogates [5]. Here, the high dimensional original system is projected onto

* Corresponding author at: The Marcus Wallenberg Laboratory for Sound and Vibration Research (MWL), Department of Engineering Mechanics, KTH Royal Institute of Technology, Stockholm, SE-100 44, Sweden.
E-mail address: rumpler@kth.se (R. Rumpler).

https://doi.org/10.1016/j.jcp.2023.112319

a lower dimensional subspace containing the desired solution while being several orders of magnitude smaller than the original solution space.

Among the projection-based methods, which are of specific interest in the present contribution, two main approaches for pMOR are typically distinguished: global and local approaches. To obtain a global basis, individual reduction matrices for some combinations of parameters are collected in global reduction matrices. As the individual bases might contain redundant information, a truncation is applied to obtain the final global reduction matrix. The size of this resulting matrix and thus the reduced-order model can grow rapidly if a large number of parameters is to be considered. In case an interpolatory method has been used to construct the individual bases, the resulting reduced model interpolates the full solution at all combinations of the parameter values which were used for the sampling process. pMOR procedures following this global approach have been presented, e.g., in [6–8]. Local pMOR methods compute reduction matrices by interpolating between the matrices obtained from other parameter combinations, respectively on a manifold of their subspaces, rather than concatenating them in a single reduction matrix [9–11]. A local method thus results in a reduced model whose size is independent of the number of local bases used. Alternatively, data-driven methods address the surrogate modelling question by computing parametric reduced-order models without requiring access to the discretisation of the original system. Such methods rely on a database of input, output, and possibly state data to compute a surrogate model depicting the behaviour of the original system. Examples for data-driven pMOR include the gappy proper orthogonal decomposition (POD) [12], the parametric Loewner framework [13], and the parametric AAA algorithm [14], among others.

In the present contribution, we present the Multivariate Well-Conditioned Asymptotic Waveform Evaluation (MWCAWE). It is a projection-based global pMOR approach extending the established Well-Conditioned Asymptotic Waveform Evaluation (WCAWE) method [15,16], which has proved to be an efficient approach for frequency sweep problems [17,18]. Contrary to many MOR approaches considering no derivatives in the parameter domain, WCAWE ensures rational interpolation up to an arbitrary order at certain expansion points. The gist of extending WCAWE to MWCAWE is to consider derivatives of the system regarding the parameters which should be retained in the reduced space. The MWCAWE basis computation involves combinations of these derivatives at each expansion point in the considered parameter space. A following recombination and truncation step ensures an appropriate size of the reduction bases corresponding to each expansion point. Multiple local bases are combined in the global pMOR approach to increase the accuracy for the complete parameter space. Although a global pMOR method may lead to larger reduced models, accuracy of the approximation and computational cost associated with the projection basis generation can justify the choice of a global method over a local approach.

In particular, the choice of parameter combinations used as expansion points for the projection basis has a major influence on the approximation quality and efficiency of the reduced-order model. A priori schemes, such as full grid, uniform, random, or Latin hypercube sampling are classical choices for systems with a small or medium number of parameters [19], but the computational cost rapidly grows with increasing dimensionality. Additionally, a priori sampling techniques are only efficient for problems with a relatively smooth parameter dependence [20]. For rougher parameter-dependence, adaptive schemes and techniques based on local sensitivity analysis may be employed to find reasonable parameter combinations as expansion points for the interpolation of the full-order model. Such methods rely on techniques estimating the approximation error of the reduced model and finding suitable locations in the parameter space for the next expansion point by, for example, a greedy algorithm; see, e.g. [21–24] and the references therein. Their associated performance relies largely on an accurate and cost-effective estimation of the approximation error, and it should be noted that these heuristic methods do not necessarily yield optimal reduced-order models. As an alternative to these sampling techniques, methods not depending on an a priori sampling of parameters have been proposed. Here, the original parameter-dependent system is reformulated to a non-parametric system where the parameter influence is modelled as additional inputs to the system [25–27]. Finally, data-driven approaches can be used to compute parametric reduced-order models; see, e.g., [28,29]. In the following, we introduce a greedy algorithm to find appropriate expansion points as a subsequent step to the MWCAWE. Thus, the original system does not have to be reformulated in any way and the approach is applicable to a wide range of dynamical systems. Contrary to most data-driven approaches MWCAWE preserves the matrix structure of the original problem and can be employed without the need for computationally intensive presampling. Targeting the bivariate case, i.e., frequency and an arbitrary parameter, the method aims at computing local reduction bases at appropriate parameter samples, such that recombinations of these bases into a global basis allow an approximation of the full-order model in the entire parametric space of interest. The algorithm relies on a combination of a contour-following algorithm and a residue-based error estimator, which is a reliable trade-off between accurately reflecting the approximation error and the involved computational cost.

The article is structured as follows: In Section 2, the univariate WCAWE and its multivariate extension, MWCAWE, are summarised. The fundamental principles of the Nested Padé Approximants, a multivariate, component-wise expansion approach based on the use of Padé Approximants are recalled in Section 3 to compare this classic approach to MWCAWE. In Section 4, the greedy multi-point scheme adopted for the sampling strategy of the bivariate case is introduced, including: the choice of the error estimator, a contour-following algorithm, and the resulting multi-patch approach. The proposed MWCAWE method and the multi-point approach detailed for the bivariate case, are then tested on two applications of increasing complexity in Section 5. A simple poro-acoustic absorption problem allows first to illustrate the potential in terms of convergence of the proposed MWCAWE, compared here to the Nested Padé Approximants on a case introduced in a pre-

vious contribution. The multi-patch, bivariate approach is then tested on a complex interior poro-acoustic problem, before concluding remarks in Section 6.

2. From the WCAWE to the MWCAWE

In this Section, the key steps from the WCAWE moment-matching approach to its multivariate extension with the MWCAWE, initially introduced in [30], are restated, in particular adapting the formalism in view of its implementation for multi-patch approximation, detailed in Section 4.

2.1. WCAWE, a projection-based model order reduction

The starting point of the WCAWE-based parametric sweep is given by a linear system,

$$\mathbf{Z}(\mathbf{x})\mathbf{u}(\mathbf{x}) = \mathbf{f}(\mathbf{x}), \tag{1}$$

where $\mathbf{x}$ is a vector of N_x independent variables corresponding to the parametric problem of interest including, e.g. the driving frequency ω, material parameters, etc. For the original univariate case discussed in [15], this vector reduces to a scalar such that $\mathbf{x} = x_1$. In a finite element (FE) problem, $\mathbf{Z}(\mathbf{x}) \in \mathbb{C}^{\mathrm{N}_n \times \mathrm{N}_n}$ represents the system matrix of the discretised problem and $\mathbf{u}(\mathbf{x}), \mathbf{f}(\mathbf{x}) \in \mathbb{C}^{\mathrm{N}_n}$ the solution vector and the vector of externally applied loads, respectively.

As N_n is typically large, the computational cost of solving (1) can be reduced by finding a system

$$\mathbf{Z}_{\mathbf{s}_n}(\mathbf{x})\boldsymbol{\alpha}_{\mathbf{s}_n}(\mathbf{x}) = \mathbf{f}_{\mathbf{s}_n}(\mathbf{x}), \tag{2}$$

of much smaller dimension $\mathrm{N_V}$ whose solution vector approximates the solution to the original system in a subspace spanned by a matrix $\mathbf{V}_{\mathbf{s}_n} \in \mathbb{C}^{\mathrm{N}_n \times \mathrm{N_V}}$, such that

$$\mathbf{u}(\mathbf{x}) \approx \hat{\mathbf{u}}_{\mathbf{s}_n}(\mathbf{x}) = \mathbf{V}_{\mathbf{s}_n}\boldsymbol{\alpha}_{\mathbf{s}_n}(\mathbf{x}). \tag{3}$$

This projection-based approach, allowing the solution of a reduced set of equations in order to obtain an approximation for (1), relies on the efficient construction of the reduction matrix $\mathbf{V}_{\mathbf{s}_n}$. A proper choice of $\mathbf{V}_{\mathbf{s}_n}$ emerges from successive derivatives of the solution vector in (1) at a specific point in the parameter space $\mathbf{x} = \mathbf{s}_n$, $\mathbf{s}_n = \left(x_{1n}, \cdots, x_{\mathrm{N}_x n}\right)$ being further referred to as the n-indexed reference point of $\mathbf{V}_{\mathbf{s}_n}$. For the original univariate case, i.e. associated with the WCAWE algorithm, $\mathrm{N}_x = 1$ and this reference point reduces to the scalar $\mathbf{s}_n = x_{1n}$. The WCAWE algorithm [15,16] provides a robust generation of this sequence of successive derivatives, overcoming the inherent ill-conditioning of the direct approach which is, e.g., used in the computation of the component-wise Padé approximants [2,18,31] and in the generation of Nested Padé approximants [3,4]. The resulting well-conditioned reduction matrix $\mathbf{V}_{\mathbf{s}_n}$, consisting of $\mathrm{N_V}$ orthonormalised basis vectors, allows for a robust, non-stagnating convergence upon increasing the size of the subspace spanned and ensures (3) in the vicinity of $\mathbf{x} = \mathbf{s}_n$. This projection

$$\mathbf{V}_{\mathbf{s}_n}^{\mathsf{H}}\mathbf{Z}(\mathbf{x})\mathbf{V}_{\mathbf{s}_n}\boldsymbol{\alpha}_{\mathbf{s}_n}(\mathbf{x}) = \mathbf{V}_{\mathbf{s}_n}^{\mathsf{H}}\mathbf{f}(\mathbf{x}), \tag{4}$$

with $\mathbf{Z}_{\mathbf{s}_n}(\mathbf{x}) = \mathbf{V}_{\mathbf{s}_n}^{\mathsf{H}}\mathbf{Z}(\mathbf{x})\mathbf{V}_{\mathbf{s}_n}$ and $\mathbf{f}_{\mathbf{s}_n}(\mathbf{x}) = \mathbf{V}_{\mathbf{s}_n}^{\mathsf{H}}\mathbf{f}(\mathbf{x})$ leads to the reduced system (2). The superscript $(\cdot)^{\mathsf{H}}$ denotes a Hermite transpose. Obviously, in the context of an FE solution, (4) is advantageously solved if the system matrix involved in the transformation, $\mathbf{Z}(\mathbf{x})$, is not reevaluated after an assembly procedure for each set of parametric vectors $\mathbf{x}$, but rather broken down in a linear combination of $\mathbf{x}$-independent global matrices.

2.2. A multivariate WCAWE basis generation procedure: the MWCAWE

A generic expression of the multivariate WCAWE algorithm (referred to as the MWCAWE in the following), as proposed in this contribution, may be given in $\mathbf{x} = \mathbf{s}_n$ by the following recursive, multiple right-hand-side (RHS) procedure adapted from the univariate case in [15,16,32],

$$
\begin{cases}
\mathbf{Z}^{(0)}\overline{\mathbf{v}}_1 = \mathbf{F}^{(0)} \\
\qquad \text{Normalisation } \overline{\mathbf{v}}_1 \longrightarrow \mathbf{v}_1 \\
\mathbf{Z}^{(0)}\overline{\mathbf{v}}_2 = \mathbf{F}^{(1)}\mathbf{e}_1^{\mathsf{T}}\mathbf{P}_{\mathbf{Q}_1}(2,1)\mathbf{e}_1 - \mathbf{Z}^{(1)}\mathbf{v}_1 \\
\qquad \text{Orthonormalisation } \overline{\mathbf{v}}_2 \longrightarrow \mathbf{v}_2 \\
\qquad \vdots \\
\mathbf{Z}^{(0)}\overline{\mathbf{v}}_n = \left(\sum_{j=1}^{(n-1)} \left(\mathbf{F}^{(j)}\mathbf{e}_1^{\mathsf{T}}\mathbf{P}_{\mathbf{Q}_1}(n,j)\mathbf{e}_{n-j} \right) - \mathbf{Z}^{(1)}\mathbf{v}_{n-1} \right. \\
\qquad\qquad \left. - \sum_{j=2}^{(n-1)} \left(\mathbf{Z}^{(j)}\mathbf{v}_{n-j}\mathbf{P}_{\mathbf{Q}_2}(n,j)\mathbf{e}_{n-j} \right) \right) \\
\qquad \text{Orthonormalisation } \overline{\mathbf{v}}_n \longrightarrow \mathbf{v}_n \\
\qquad \vdots \\
\mathbf{Z}^{(0)}\overline{\mathbf{v}}_{\mathrm{N_V}} = \left(\sum_{j=1}^{(\mathrm{N_V}-1)} \left(\mathbf{F}^{(j)}\mathbf{e}_1^{\mathsf{T}}\mathbf{P}_{\mathbf{Q}_1}(\mathrm{N_V},j)\mathbf{e}_{\mathrm{N_V}-j} \right) - \mathbf{Z}^{(1)}\mathbf{v}_{\mathrm{N_V}-1} \right. \\
\qquad\qquad \left. - \sum_{j=2}^{(\mathrm{N_V}-1)} \left(\mathbf{Z}^{(j)}\mathbf{v}_{\mathrm{N_V}-j}\mathbf{P}_{\mathbf{Q}_2}(\mathrm{N_V},j)\mathbf{e}_{\mathrm{N_V}-j} \right) \right) \\
\qquad \text{Orthonormalisation } \overline{\mathbf{v}}_{\mathrm{N_V}} \longrightarrow \mathbf{v}_{\mathrm{N_V}}
\end{cases}
\tag{5}
$$

where a modified Gram-Schmidt orthonormalisation step is performed between each vector generation by the multiple RHS systems in (5). Furthermore

- $\mathbf{Z}^{(k)}$ stands for the differentiation of the system matrix $\mathbf{Z}$, in $\mathbf{x} = \mathbf{s}_n$, to the "cumulative" multivariate order k, regardless of the distribution of these partial derivatives between the independent variables x_j, $j = 1, \ldots, \mathrm{N}_x$. In other words, given the notation $\partial_{x_j}^{i_j}(\cdot)$ for the partial derivative with respect to x_j at order i_j, and the convention $\partial_{x_j}^{0}(\cdot) = (\cdot)$, then $\mathbf{Z}^{(k)}$ results from the summed partial derivative orders such that

$$
\mathbf{Z}^{\left(\sum_{j=1}^{\mathrm{N}_x} i_j\right)} = \left(\prod_{j=1}^{\mathrm{N}_x} \partial_{x_j}^{i_j} \right) \mathbf{Z}, \tag{6}
$$

  with the aforementioned convention implying that $\mathbf{Z}^{(0)} = \mathbf{Z}$ (dependency on $\mathbf{x}$ omitted).
- $\mathbf{e}_k$ is a unitary standard basis vector associated with the k^{th} component of the solution vector.
- $\overline{\mathbf{v}}_k$ is the non-orthonormalised vector generated in the k^{th} iteration of the procedure.
- $\mathbf{v}_k$ is the basis vector orthonormalised against $\mathbf{v}_1, \ldots, \mathbf{v}_{k-1}$, generated after the k^{th} iteration of the procedure.
- $\mathbf{P}_{\mathbf{Q}_\omega}(\alpha, \beta)$, $\omega = 1, 2$, corresponds to the RHS correction terms, chosen to be associated with the modified Gram-Schmidt orthonormalisation process [15].

The resulting orthonormalised and non-orthonormalised bases, $\mathbf{V}_{\mathbf{s}_n} = \left[\mathbf{v}_1 \cdots \mathbf{v}_{\mathrm{N_V}}\right]$ and $\overline{\mathbf{V}}_{\mathbf{s}_n} = \left[\overline{\mathbf{v}}_1 \cdots \overline{\mathbf{v}}_{\mathrm{N_V}}\right]$ respectively, are related by

$$
\mathbf{V}_{\mathbf{s}_n} = \overline{\mathbf{V}}_{\mathbf{s}_n}\mathbf{Q}^{-1}, \tag{7}
$$

where $\mathbf{Q}$ is an $\mathrm{N_V} \times \mathrm{N_V}$ upper triangular, non-singular matrix containing the modified Gram-Schmidt coefficients. More precisely, column k of $\mathbf{Q}$ contains the successive coefficients resulting from the projection of the partially orthonormalised $\overline{\mathbf{v}}_k$ on the orthonormalised vectors $\mathbf{v}_j$, $j < k$, and Q_{kk} corresponds to the norm of $\mathbf{v}_k$ before its normalisation. The correction terms $\mathbf{P}_{\mathbf{Q}_\omega}(\alpha, \beta)$, $\omega = 1, 2$, are given by the following product of block matrices extracted from $\mathbf{Q}$,

$$
\mathbf{P}_{\mathbf{Q}_\omega}(\alpha, \beta) = \prod_{t=\omega}^{\beta} \mathbf{Q}^{-1}_{[t:\alpha-\beta+t-1, t:\alpha-\beta+t-1]}. \tag{8}
$$

Further discussions on the choice of the RHS correction coefficients other than associated with the Gram-Schmidt coefficients, may be found in [15]. The approximated solution may be evaluated at all degrees of freedom (DOFs) from the generalised coordinates vector $\boldsymbol{\alpha}_{\mathbf{s}_n}(\mathbf{x})$ following (3).

If during the sequence (5) the orthonormalised vector $\mathbf{v}_n$ would be linearly dependent to the previous vector $\mathbf{v}_{n-1}$, numerical breakdown occurs. In this case the algorithm stops prematurely, similar to the closely related Arnoldi processes [33,34]. A breakdown in this setting implies that the subspace already contains the information which would be added by computing further moments. However, the orthogonality of the basis vectors is imposed by the correction terms $\mathbf{P}_{\mathbf{Q}_1}$ and $\mathbf{P}_{\mathbf{Q}_2}$ in (5) and the basis vectors are thus well-conditioned [35,36].

Adapting the WCAWE algorithm from univariate to multivariate problems involves the generation of sequences of RHS vectors emerging from sequences of iteratively differentiated matrices $\mathbf{Z}^{(k)}$. The choice adopted here is to generate such sequences independently, thus generating a set of N_Σ bases $\left\{\mathbf{V}^1_{\mathbf{s}_n}, \cdots, \mathbf{V}^{\mathrm{N}_\Sigma}_{\mathbf{s}_n}\right\}$ associated with N_Σ sequences of iteratively differentiated matrices $\mathbf{Z}^{(k)}$. This procedure is illustrated in Fig. 1. These bases, in principle all of the same size $\mathrm{N_V}$,[1] each consist of orthonormalised basis vectors. However, these are naturally not mutually orthonormalised from one basis to the next. Additionally, it is not guaranteed that the N_Σ-sequences produce linearly independent subsets of vectors upon merging the resulting bases. A simple way to ensure this, as well as to reduce the basis to a minimum number of basis vectors, is to proceed with a compression, or a component selection, via a Singular Value Decomposition (SVD) of the merged set of bases [37,38]. Either way, the initial step consists in an SVD of the merged basis after concatenation, $\mathbf{V}^{\mathrm{mer}}_{\mathbf{s}_n} = \left[\mathbf{V}^1_{\mathbf{s}_n} \cdots \mathbf{V}^{\mathrm{N}_\Sigma}_{\mathbf{s}_n}\right]$, resulting in

$$\mathbf{V}^{\mathrm{mer}}_{\mathbf{s}_n} = \sum_{i=1}^{\mathrm{N_{V^{mer}}}} \sigma_i \mathbf{w}^{\mathrm{l}}_i \mathbf{w}^{\mathrm{r}\mathsf{T}}_i, \tag{9}$$

where $\mathrm{N_{V^{mer}}} \leqslant \mathrm{N_V N}_\Sigma$ corresponds to the total number of basis vectors in $\mathbf{V}^{\mathrm{mer}}_{\mathbf{s}_n}$ and σ_i, $\mathbf{w}^{\mathrm{l}}_i$, and $\mathbf{w}^{\mathrm{r}}_i$ correspond to the singular values and left and right singular vectors, respectively. $\mathbf{V}^{\mathrm{mer}}_{\mathbf{s}_n}$ is expected to have a moderate number of columns, such that this additional SVD may be performed with relatively low computational effort compared to the matrix decompositions required in the WCAWE.

From the SVD, the (left) singular vectors associated with the highest singular values are selected as the components of the final local reduction matrix, such that assuming a descending-ordered sequence of singular values $\left(\sigma_1 \cdots \sigma_{\mathrm{N_{V^{mer}}}}\right)$, the reduced, merged basis $\mathbf{V}^*_{\mathbf{s}_n}$ is

$$\mathbf{V}^*_{\mathbf{s}_n} = \left\{\mathbf{w}^{\mathrm{l}}_i \mid i \in \{1, \ldots, \mathrm{N_{V^{mer}}}\} \wedge \sigma_i \geqslant \sigma_{\mathrm{thresh}}\right\}, \tag{10}$$

where σ_{thresh} corresponds to the empirically chosen threshold value for the selection of the singular values resulting from the decomposition in (9).

The resulting local projection basis $\mathbf{V}^*_{\mathbf{s}_n}$, consisting of $\mathrm{N_{V^*}}$ vectors, is then introduced in place of the generic, univariate transformation corresponding to (3), such that

$$\hat{\mathbf{u}}^*_{\mathbf{s}_n}(\mathbf{x}) = \mathbf{V}^*_{\mathbf{s}_n}\, \boldsymbol{\alpha}^*_{\mathbf{s}_n}(\mathbf{x}), \tag{11}$$

leading to a reduced system associated with the initial problem in (1), such that

$$\left(\mathbf{V}^*_{\mathbf{s}_n}\right)^{\mathsf{H}} \mathbf{Z}(\mathbf{x}) \mathbf{V}^*_{\mathbf{s}_n} \boldsymbol{\alpha}^*_{\mathbf{s}_n}(\mathbf{x}) = \left(\mathbf{V}^*_{\mathbf{s}_n}\right)^{\mathsf{H}} \mathbf{f}(\mathbf{x}), \tag{12}$$

involving $\mathrm{N_{V^*}} \leqslant \mathrm{N_{V^{mer}}}$ generalised coordinates in $\boldsymbol{\alpha}^*_{\mathbf{s}_n}(\mathbf{x})$.

2.3. Sequences of multivariate differentiation

The MWCAWE procedure above depends on the generation of sequences of differentiated system matrices to successive "cumulated" multivariate partial derivative orders k, denoted $\mathbf{Z}^{(k)}$, $k = 0, \ldots, \mathrm{N_V} - 1$. The choice of these sequences is addressed empirically in this contribution, and illustrated for the case of bivariate problems. For the bivariate case, provided orders of partial derivatives N_{x_1} and N_{x_2} with respect to x_1 and x_2, respectively, such that $\mathrm{N}_{x_1} + \mathrm{N}_{x_2} = \mathrm{N_V} - 1$, the resulting cumulated order of partial derivatives and the associated differentiation paths may be represented in a matrix form, see Fig. 1. The matrix sequences obtained from these differentiation paths of the system matrix subsequently feed the procedure presented in (5), ultimately resulting in the sequence of bases concatenated into $\mathbf{V}^{\mathrm{mer}}_{\mathbf{s}_n}$, see (9).

Rather than computing all possible differentiation paths for a given set of maximum orders N_{x_1} and N_{x_2}, only a subset of these is computed. Preliminary tests based on the methodology introduced in the previous section rapidly confirm the emergence of quasi-linearly dependent basis vectors associated with different partial derivative orders from distinct sequences. In practice, it is estimated that a number of sequences of the order of the dimensionality of the multivariate

[1] Strictly speaking, the size of the bases in this N_Σ-sequence decreases by a dimensionality-dependent number of vectors for each basis added to the sequence, due to the partial overlap of consecutive sequences, see, e.g., Section 2.3 and Fig. 1. This point is however marginally relevant here given the SVD step involved in the subsequent processing of these bases, see (9) and (10).

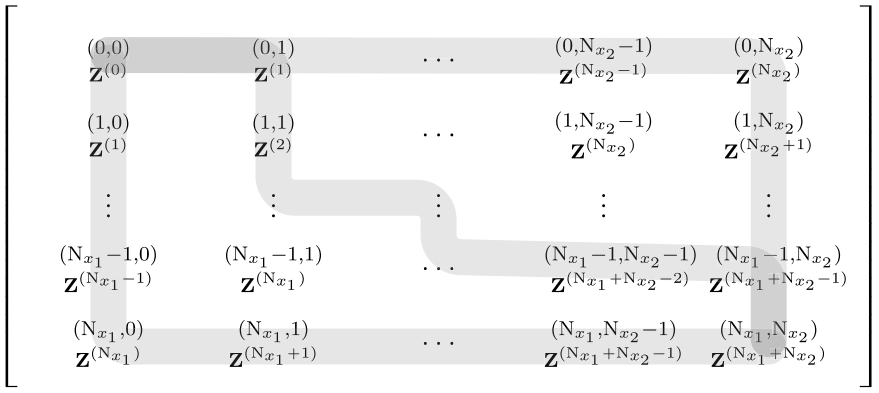

Fig. 1. Illustration of three possible bivariate derivation paths for the system matrix $\mathbf{Z}$ up to the "cumulated" order of partial derivatives $(\mathrm{N}_{x_1}+\mathrm{N}_{x_2})$.

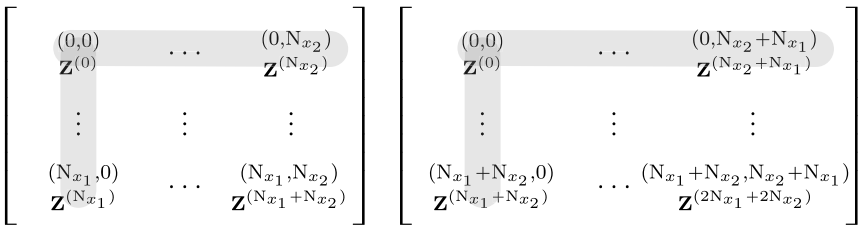

Fig. 2. Illustration of alternative derivation paths involving univariate sequences.

problem may be sufficient in order to provide good convergence properties of the ROM, i.e., $\mathrm{N}_\Sigma = \mathrm{N}_x$. For instance, $\mathrm{N}_\Sigma = 2$, i.e. two sequences of differentiation, may be sufficient for a bivariate problem. This point was preliminarily tested in [30], where a comparison between merging bases from a two-sequence and a three-sequence approach was reported, with the following observations:

- When introducing a three-sequence approach, only a few of the vectors are discarded by the SVD step in (10) for a threshold $\sigma_{\text{thresh}} = 10^{-15} \cdot \sigma_{\max}$, where $\sigma_{\max}$ corresponds to the largest singular value in the decomposition of (10);
- Conversely, a two-sequence approach for the same truncation criterion, and based on the two outermost sequences in Fig. 1, was observed not to discard any of the singular values and associated vectors;
- A comparison between the error for the approximation resulting from a two-sequence and a three-sequence approach showed no significant improvement of convergence and accuracy, thus indicating only minor contribution from the basis vectors emerging from the additional sequence.

Note that this result addresses in part the issue of dimensionality and the associated efficiency of the approach: the number of necessary successive partial derivatives grows linearly with the dimensionality as opposed to an exponential growth if all cross-derivatives are calculated in a fully multivariate approach.

Among the several alternative derivation paths involving two sequences, two special cases may be interesting to highlight in view of the choice of truncation of these sequences. Fig. 2 illustrates configurations where, based on the previously introduced orders N_{x_1} and N_{x_2}, only univariate derivation sequences are considered: first considering shorter derivation paths where the cross derivatives are removed from the paths in Fig. 1, obviously leading to a reduced set of $\mathrm{N}_{\mathrm{V}^{\text{mer}}} = \mathrm{N}_{x_1} + \mathrm{N}_{x_2} + 1$ basis vectors, see Fig. 2a; second, keeping a similar size of reduction basis to the one associated with the suggestion in Fig. 1, the "cumulated" order of partial derivatives $(\mathrm{N}_{x_1} + \mathrm{N}_{x_2})$ is reallocated to sequences of univariate derivatives, see Fig. 2b. These three alternatives are considered further in the application Section 5.

3. A comparison to Padé-based expansion methods: the Nested Padé Approximants

In a previous contribution of one of the authors, a multivariate rational expansion method based on Padé Approximants was implemented in order to address problems of the form of (1), together with a multi-patch strategy for bivariate problems [4]. This approach is briefly outlined in the following to highlight the convergence improvements offered by an approach such as the MWCAWE, as further detailed in the applications, see Section 5. Although the Nested Padé Approximants have the benefit of providing piecewise analytical expressions of the solution for each DOF, the ill-conditioned nature of the associated iterative process leads to the same limitations in convergence as for the underlying univariate case. This

contrast of convergence between the component-wise Padé-based univariate expansion and the WCAWE was, for example, evaluated in [18].

In concise terms, the approximation of a multivariate problem, (1), may be performed by expanding the solution for each DOF in the form of nested univariate Padé Approximants. For a bivariate problem where $\mathbf{x} = (x_1, x_2)$, the expansion $\hat{u}(\Delta x_1, \Delta x_2)$ of a component $u(x_1, x_2)$ of the solution vector $\mathbf{u}(x_1, x_2)$ around the reference point $\mathbf{s}_n = (x_{1n}, x_{2n})$, is given by

$$\hat{u}(\Delta x_1, \Delta x_2) = \frac{\displaystyle\sum_{i=0}^{\mathrm{m}_1} \frac{\displaystyle\sum_{j=0}^{\mathrm{m}_2} t_{ij} \Delta {x_2}^j}{1 + \displaystyle\sum_{j=1}^{\mathrm{n}_2} t_{i(\mathrm{m}_2+j)} \Delta {x_2}^j} \Delta {x_1}^i}{1 + \displaystyle\sum_{i=1}^{\mathrm{n}_1} \frac{\displaystyle\sum_{j=0}^{\mathrm{m}_2} t_{(\mathrm{m}_1+i)j} \Delta {x_2}^j}{1 + \displaystyle\sum_{j=1}^{\mathrm{n}_2} t_{(\mathrm{m}_1+i)(\mathrm{m}_2+j)} \Delta {x_2}^j} \Delta {x_1}^i}, \tag{13}$$

where $\begin{bmatrix} \Delta x_1 & \Delta x_2 \end{bmatrix} = \begin{bmatrix} x_1 & x_2 \end{bmatrix} - \begin{bmatrix} x_{1n} & x_{2n} \end{bmatrix}$; m_k and n_k, $k = 1, 2$, correspond to the orders of expansion, distributed between the numerator and denominator of the nested Padé Approximants, where $\mathrm{N}_{xk} = \mathrm{m}_k + \mathrm{n}_k$; t_{ij} are the $\left(\mathrm{N}_{x_1} + 1\right)\left(\mathrm{N}_{x_2} + 1\right)$ nested Padé coefficients determined recursively from linear systems of equations obtained from a Taylor expansion of $u(x_1, x_2)$ around the reference point $\mathbf{s}_n = (x_{1n}, x_{2n})$; see [3,4] for further details. The extension to higher dimensionality may be illustrated, e.g., for $\mathbf{x} = (x_1, x_2, x_3)$, adopting similar conventions and notations, as

$$\hat{u}(\Delta x_1, \Delta x_2, \Delta x_3) = \frac{P(\Delta x_1, \Delta x_2, \Delta x_3)}{Q(\Delta x_1, \Delta x_2, \Delta x_3)} \tag{14}$$

with

$$P(\Delta x_1, \Delta x_2, \Delta x_3) = \sum_{i=0}^{\mathrm{m}_1} \frac{\displaystyle\sum_{j=0}^{\mathrm{m}_2} \frac{\displaystyle\sum_{k=0}^{\mathrm{m}_3} t_{ijk} \Delta {x_3}^k}{1 + \displaystyle\sum_{k=1}^{\mathrm{n}_3} t_{ij(\mathrm{m}_3+k)} \Delta {x_3}^k} \Delta {x_2}^j}{1 + \displaystyle\sum_{j=1}^{\mathrm{n}_2} \frac{\displaystyle\sum_{k=0}^{\mathrm{m}_3} t_{i(\mathrm{m}_2+j)k} \Delta {x_3}^k}{1 + \displaystyle\sum_{k=1}^{\mathrm{n}_3} t_{i(\mathrm{m}_2+j)(\mathrm{m}_3+k)} \Delta {x_3}^k} \Delta {x_2}^j} \Delta {x_1}^i,$$

$$Q(\Delta x_1, \Delta x_2, \Delta x_3) = 1 + \sum_{i=1}^{\mathrm{n}_1} \frac{\displaystyle\sum_{j=0}^{\mathrm{m}_2} \frac{\displaystyle\sum_{k=0}^{\mathrm{m}_3} t_{(\mathrm{m}_1+i)jk} \Delta {x_3}^k}{1 + \displaystyle\sum_{k=1}^{\mathrm{n}_3} t_{(\mathrm{m}_1+i)j(\mathrm{m}_3+k)} \Delta {x_3}^k} \Delta {x_2}^j}{1 + \displaystyle\sum_{j=1}^{\mathrm{n}_2} \frac{\displaystyle\sum_{k=0}^{\mathrm{m}_3} t_{(\mathrm{m}_1+i)(\mathrm{m}_2+j)k} \Delta {x_3}^k}{1 + \displaystyle\sum_{k=1}^{\mathrm{n}_3} t_{(\mathrm{m}_1+i)(\mathrm{m}_2+j)(\mathrm{m}_3+k)} \Delta {x_3}^k} \Delta {x_2}^j} \Delta {x_1}^i,$$

highlighting the nested nature of the expansion. This structure is the primary limiting factor for applying the approach to problems with higher dimensional parameter dependence as the associated recursive procedure propagates numerical errors due to ill-conditioned systems of equations.

4. Multi-patch, MWCAWE approach

In practice, a reduced-order model computed from a single reference point $\mathbf{s}_n$ may not be able to efficiently approximate the full-order model in the desired range of parameters. Introducing multiple expansion points $\mathbf{s}_n$ potentially increases the accuracy of the reduced-order model for wider parameter ranges, but each new expansion point requires the decomposition of the full order system matrix. Increasing the order of approximation is typically less computationally expensive, but the approximation quality may not significantly benefit from the higher orders and eventually stagnate. This implies a trade-off between increasing the order of approximation, reflected in the size of the basis $\mathrm{N_V}$ for a given reference point $\mathbf{s}_n$, and a multi-patch strategy involving several such reference points. In contrast, a multi-patch approximation strategy was mostly driven by convergence limitations for the Nested Padé Approximants approach [4].

Regardless of the motivation for a multi-patch approximation, an assessment of the approximation error is a necessary prior step for an efficient, automated sampling of the parametric space, and discussed in the following section.

4.1. Error estimator

A residue-based error estimator is used to identify bounds of the patches of the approximated solution with a satisfying degree of accuracy. Such residual-based estimators have shown to be both efficient and accurate ways of tracking the approximation error [4,39–45]. Several contributions in the past have chosen a Cauchy-type of convergence, where two or more approximations of different orders or with different expansion points are evaluated and compared [19,46–48]. Patches of acceptable approximation are subsequently defined within the bounds where these solutions agree to a satisfying degree of accuracy. Although this approach has proved to be both easy to implement and reasonably accurate, it entails limitations in the scope of the multi-patch, multivariate approach proposed here.

First, this Cauchy-type of convergence is particularly convenient for monotonic sequences, as a reasonably good correlation may be found between the bounds generated from the comparison of two sequential approximations and those associated with the actual approximation error. However, the WCAWE was highlighted, e.g. in [18], to have an occasional, punctual increase of the error in the convergence sequence, or stagnation of the error. Such occurrences may require a special implementation of the error estimation, e.g., relying on more than two sequential evaluations. Although the computational cost associated with these reduced models is in principle only a fraction of the cost associated with the full-sized model, the potential need to evaluate several sequential approximations may become increasingly impractical as dimensionality increases.

Second, beyond the cost associated with the evaluation of several approximations, such an approach may not align naturally with the need to limit the number of evaluations of the approximation error to be performed in the context of multivariate problems. In fact, for univariate or low-dimensional problems, the error is in most cases evaluated in all points of the parametric sweep. But doing so for multivariate problems is both unnecessary and rapidly overly time-consuming.

In view of these conditions, the residue-based error estimator approach proposed here seeks to be relying on a unique set of local bases merged into $\mathbf{V}^*_{\mathbf{s}_n}$, see (10), and a single evaluation of the error limited to a subset of necessary points in the parametric space. The residue is evaluated on the basis of the residual vector $\mathbf{r}^*_{\mathbf{s}_n}(\mathbf{x})$ associated with the approximated solution $\hat{\mathbf{u}}^*_{\mathbf{s}_n}(\mathbf{x})$, see (11) and (12), following a transformation with the merged local bases $\mathbf{V}^*_{\mathbf{s}_n}$ of (10). This residual vector is given by

$$\mathbf{r}^*_{\mathbf{s}_n}(\mathbf{x}) = \mathbf{Z}(\mathbf{x})\,\hat{\mathbf{u}}^*_{\mathbf{s}_n}(\mathbf{x}) - \mathbf{f}(\mathbf{x}), \tag{15}$$

for a generic problem of the form of (1). It may subsequently be used as an error indicator based on its ℓ^2 norm, such that,

$$\varepsilon^*_{\mathbf{s}_n}(\mathbf{x}) = \frac{||\mathbf{r}^*_{\mathbf{s}_n}(\mathbf{x})||_2}{\min\limits_{\mathbf{y}\in\Omega\setminus\{\mathbf{s}_n\}}\left(||\mathbf{r}^*_{\mathbf{s}_n}(\mathbf{y})||_2\right)}, \tag{16}$$

which is normalised with respect to the minimum residual norm evaluated, assuming Ω to be a set of all vectors $\mathbf{x}$ in the parametric space to be evaluated. This minimum residual norm is thus logically corresponding to an $\mathbf{x}$ in the closest vicinity of $\mathbf{s}_n$, given that an evaluation in $\mathbf{s}_n$ is excluded. This error estimator enables the search for bounds of convergence associated with each pair of basis and reference point, as further detailed in the next section.

4.2. The bivariate case: contour-following algorithm

In order to establish a multi-patch strategy, i.e. defining a set of points $\{\mathbf{s}_1,\ldots,\mathbf{s}_n,\ldots,\mathbf{s}_\mathrm{N}\}$ at which local bases are to be calculated, the bounds of accuracy associated with each of these bases have to be estimated. This is here proposed to be done assuming a monotonic or at least smooth increase of the approximation error as the distance to the reference parametric point $\mathbf{s}_n$ increases. This assumption, in principle verified given the limit of relatively small approximation errors, as is the case for the considered applications where a high degree of accuracy is sought, leads to seeking patches of convergence in the form of simply-connected domains. A Moore-Neighbour contour tracing algorithm is implemented for this purpose, which follows an error estimator isocontour corresponding to the predetermined tolerance $\varepsilon_{\max}$. Note that

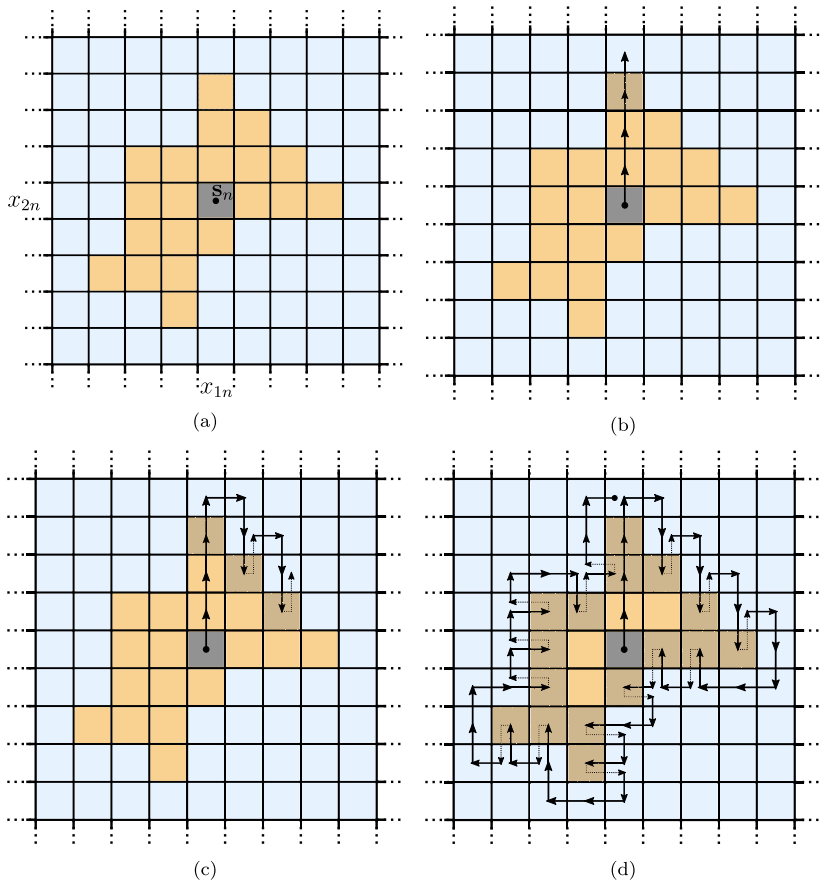

Fig. 3. Illustration of the Moore-Neighbour contour tracing algorithm for a single patch. Light orange: domain to be uncovered where $\varepsilon^*_{\mathbf{s}_n}(\mathbf{x}) \leqslant \varepsilon_{\max}$; light blue: domain outside the contour to be uncovered where $\varepsilon^*_{\mathbf{s}_n}(\mathbf{x}) > \varepsilon_{\max}$; dark grey: arbitrary reference point; light grey: uncovered contour points. (a) Reference parametric point $\mathbf{s}_n = [x_{1n}\, x_{2n}]$; (b) initial contour point in an arbitrary direction; (c) first three contour points; (d) completed Moore-Neighbour contour tracing sequence. (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)

in the case where the resulting patches would not satisfy the assumption of obtaining simply-connected domains, the inexpensive introduction of the error estimator for all the points in the domain would allow to detect such an anomaly.

The contour-following algorithm operates as follows: Starting from a parametric point among the closest to $\mathbf{s}_1 = (x_{11}, x_{21})$, for a bivariate problem, the algorithm seeks the first contour point in an arbitrarily chosen direction, e.g., in the direction of increasing x_2 for a set value of x_1 (see Figs. 3a and 3b), by estimating the error $\varepsilon^*_{\mathbf{s}_1}$ according to (16). Assuming a structured-grid point distribution such that two-dimensional square lattices may be defined for the bivariate case of interest, and thus also defining Moore neighbourhoods, the subsequent contour points are sequentially determined with the contour-following algorithm (see Figs. 3c and 3d). This approach enables the computation of the error estimator at a reduced number of points only, which lie closest to the estimated bound of convergence for the reduced system associated with the basis $\mathbf{V}^*_{\mathbf{s}_1}$.

The estimation of such bounds of convergence subsequently allows to choose sequentially locations for new reference points $\mathbf{s}_n = (x_{1n}, x_{2n})$, $n = 2, \ldots, \mathrm{N}$, for the calculation of additional local bases. This opens the way for a multi-patch solution strategy, gradually filling the entire bivariate range of interest with sub-domains of satisfactory approximation. An algorithm computing automatically such a multi-patch solution is presented in the following.

4.3. The bivariate case: a multi-patch algorithm

As introduced in connection with (16), let Ω be a set of all vectors $\mathbf{x}$ where the solution to (1) is sought to be approximated, and $\partial\Omega$ its outer boundary set. A sequential approach is considered in order to establish N local merged bases associated with a sequence of N reference points $\langle \mathbf{s}_n | \mathbf{s}_n \in \Omega \wedge n \in \{1, \ldots, \mathrm{N}\}\rangle$. It aims at gradually filling the parametric domain of interest with converged sub-patches. Each patch is bounded by a contour where the accuracy of the reduced-order model associated with the corresponding reference point is estimated to be sufficient. The number of local merged bases necessary to approximate the solution in the complete domain of interest Ω is iteratively determined given a tolerance for the error estimator $\varepsilon_{\max}$, the bounds of the parametric space of interest $\partial\Omega$, as well as a gap-tolerance $\Delta_{\mathrm{gap\_max}}$, in

terms of the maximum allowable distance between two neighbour domains of convergence. In this region none of the local merged bases $\mathbf{V}^*_{\mathbf{s}_n}$ fulfils the convergence criterion $\varepsilon^*_{\mathbf{s}_n}(\mathbf{x}) \leqslant \varepsilon_{\max}$. Since the introduction of a residual-based error estimator alleviates the need for multiple reduced-order models resulting from complementary approximations of the original model, the dimensions of the reduction bases are fixed a priori and set to be equal for all local bases.

Further, the set Ω_{conv} is introduced, comprising the parameter combinations for which at least one of the merged local bases fulfils the convergence criterion,

$$\Omega_{\text{conv}} = \left\{ \mathbf{x} \mid \varepsilon^*_{\mathbf{s}_n}(\mathbf{x}) \leqslant \varepsilon_{\max} \wedge n \in \{1, \dots, \mathrm{N}\} \right\}. \tag{17}$$

$\partial\Omega_{\text{conv}}$ is subsequently introduced as its boundary set, while its relative complement, where no local basis allows to reach an acceptable approximation, is denoted $\Omega_{\text{gap}} = \Omega \setminus \Omega_{\text{conv}}$. Additionally, associated with each of the reference points $\mathbf{s}_n$, the local merged basis denoted $\mathbf{V}^*_{\mathbf{s}_n}$ allows to approximate the solution for a subset $\Omega_{\mathbf{s}_n} \subseteq \Omega_{\text{conv}}$. This local basis $\mathbf{V}^*_{\mathbf{s}_n}$ is obtained according to the details introduced in Section 2.2, leading to (10), whose procedure is referred to in Algorithm 1 as the function GetBasis(). For each local basis, the isocontour associated with the error estimator tolerance $\varepsilon_{\max}$ is obtained according to the details of the contour-following procedure introduced in Section 4.2, and referred to as the function GetContour(). The closure associated with the boundary set $\partial\Omega_{\mathbf{s}_n}$ is obtained via a function referred to as Fill(). Finally, the function d() refers to the signed distance function, returning the distance of a parameter combination $\mathbf{x}$ to a boundary $\partial\Omega$ or $\partial\Omega_{\text{conv}}$, assuming here a convention of negative values inside Ω_{conv} and positive outside. A conceptual pseudocode of the multi-patch generation of these N local bases is presented in Algorithm 1.

Algorithm 1: Calculation of the multi-patch sub-bases.

Result: A set of sub-bases with their estimated bounds of convergence

```
// Initialisation
ε_max ← Max error tolerance
Δ_gap_max ← Max allowable gap distance to any patch
Define Ω, ∂Ω, Ω_conv = ∅                                   // Domain and boundary sets
Ω_gap ← Ω \ Ω_conv
// Starting basis from initial reference point s_1
s_1 ← (x_11 x_21)
V*_s1 ← GetBasis(s_1, N_V, N_Σ)
∂Ω_s1 ← GetContour(s_1, V*_s1, ε_max)
Ω_s1 ← Fill(∂Ω_s1)
// Update convergence and gap sets
Ω_conv ← Ω_conv ∪ Ω_s1
Ω_gap ← Ω \ Ω_conv
Gap ← Max {d(x, ∂Ω_conv ∪ ∂Ω) | x ∈ Ω_gap}
n ← 1
// Calculation of subsequent sub-bases
while Gap > Δ_gap_max do
    n ← n + 1
    s_n ← Argmax {d(x, ∂Ω_conv ∪ ∂Ω) | x ∈ Ω_gap}
    V*_sn ← GetBasis(s_n, N_V, N_Σ)
    ∂Ω_sn ← GetContour(s_n, V*_sn, ε_max)
    Ω_sn ← Fill(∂Ω_sn)
    // Update convergence and gap sets
    Ω_conv ← Ω_conv ∪ Ω_sn
    Ω_gap ← Ω \ Ω_conv
    Gap ← Max {d(x, ∂Ω_conv ∪ ∂Ω) | x ∈ Ω_gap}
end
```

Algorithm 1 leads to a set of N local bases $\mathbf{V}^*_{\mathbf{s}_n}$ associated with the reference points $\mathbf{s}_n$, $n = 1, \dots, \mathrm{N}$. It is illustrated for a case consisting of four reference points in Fig. 4. Depending on the choice of tolerance $\Delta_{\text{gap\_max}}$, the domains of convergence associated with these local bases are not necessarily overlapping, which may lead to gaps in the domain where the approximation of the original solution is not fulfilling the sought degree of accuracy. It is, however, noteworthy that by construction, these bases are not expected to be consisting of vectors mutually linearly independent. According to the nature of the dependence of the response with respect to the set of variables in $\mathbf{x}$, neighbouring local bases may contribute to enriching the subspace spanned by each basis $\mathbf{V}^*_{\mathbf{s}_n}$, $n = 1, \dots, \mathrm{N}$. These local bases may thus be advantageously recombined into a global basis, in particular in order to improve the accuracy of the approximation in the gaps, where no specific local basis is initially assigned. This recombination is in the following examples applied to form a single global basis, which may however be easily adapted to only subsets of neighbouring local bases in cases of very large parametric domains. A natural way to recombine these is via an SVD, as already introduced for the sequences of local bases emerging from the partial differentiation process in Section 2.2, implicitly performed in function GetBasis() of Algorithm 1. Similarly to (9), the merged basis matrix $\mathbf{V}^{*\text{mer}} = \left[\mathbf{V}^*_{\mathbf{s}_1} \cdots \mathbf{V}^*_{\mathbf{s}_\mathrm{N}}\right]$ is decomposed following an SVD such that

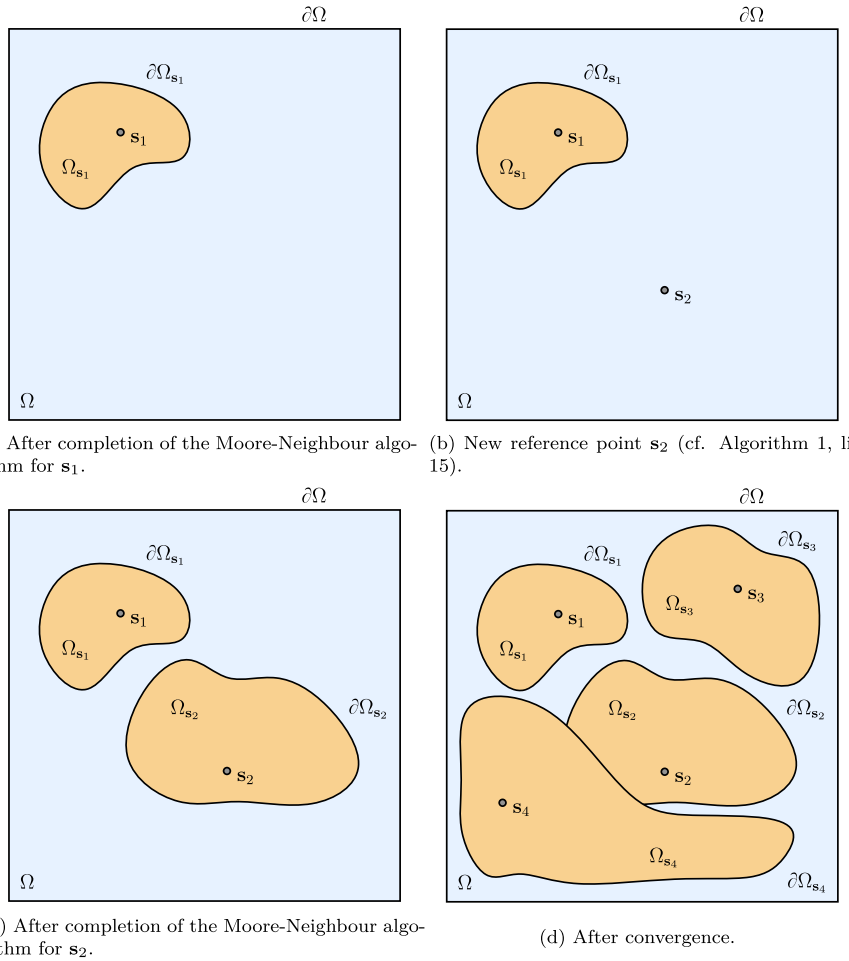

(a) After completion of the Moore-Neighbour algorithm for $\mathbf{s}_1$.

(b) New reference point $\mathbf{s}_2$ (cf. Algorithm 1, line 15).

(c) After completion of the Moore-Neighbour algorithm for $\mathbf{s}_2$.

(d) After convergence.

Fig. 4. Illustration of Algorithm 1. The domain of interest Ω is depicted in light blue, the converged contours $\Omega_{\mathbf{s}_n}$ in light orange.

$$\mathbf{V}^{*\mathrm{mer}} = \sum_{i=1}^{\mathrm{N}_{\mathbf{V}^{*\mathrm{mer}}}} \sigma_i^* \mathbf{w}_i^{\mathrm{l}*} \mathbf{w}_i^{\mathrm{r}*\mathsf{T}}, \tag{18}$$

where $\mathrm{N}_{\mathbf{V}^{*\mathrm{mer}}}$ corresponds to the total number of basis vectors in $\mathbf{V}^{*\mathrm{mer}}$; σ_i^*, $\mathbf{w}_i^{\mathrm{l}*}$, and $\mathbf{w}_i^{\mathrm{r}*}$ correspond to the singular values and left and right singular vectors respectively. From the SVD, the left singular vectors associated with the highest singular values are selected as the components of the reduced merged basis $\mathbf{V}^{**}$, such that assuming a descending-ordered sequence of singular values $\left[\sigma_1^* \cdots \sigma_{\mathrm{N}_{\mathbf{V}^{*\mathrm{mer}}}}^*\right]$, the reduced, merged basis is given by

$$\mathbf{V}^{**} = \left\{ \mathbf{w}_i^{\mathrm{l}*} \mid i \in \{1, \ldots, \mathrm{N}_{\mathbf{V}^{*\mathrm{mer}}}\} \wedge \sigma_i^* \geqslant \sigma_{\mathrm{thresh}} \right\}, \tag{19}$$

where σ_{thresh} corresponds to an empirically chosen threshold value for the selection of the singular values resulting from the decomposition in (18).

The resulting global reduction matrix $\mathbf{V}^{**}$, consisting of $\mathrm{N}_{\mathbf{V}^{**}}$ vectors, is then introduced in place of the single-patch transformation corresponding to (11), such that,

$$\hat{\mathbf{u}}^{**}(\mathbf{x}) = \mathbf{V}^{**} \boldsymbol{\alpha}^{**}(\mathbf{x}). \tag{20}$$

This leads to a reduced system associated with the initial problem in (1), such that,

$$\left(\mathbf{V}^{**}\right)^{\mathsf{T}} \mathbf{Z}(\mathbf{x}) \mathbf{V}^{**} \boldsymbol{\alpha}^{**}(\mathbf{x}) = \left(\mathbf{V}^{**}\right)^{\mathsf{T}} \mathbf{f}(\mathbf{x}), \tag{21}$$

involving $\mathrm{N}_{\mathbf{V}^{**}} \leqslant \mathrm{N}_{\mathbf{V}^{*\mathrm{mer}}}$ generalised coordinates in $\boldsymbol{\alpha}^{**}(\mathbf{x})$.

It is worth noting that Algorithm 1 may easily be modified to be applicable to problems with more than two parameters: The underlying model order reduction via MWCAWE is applicable to problems with an arbitrary number of parameters and

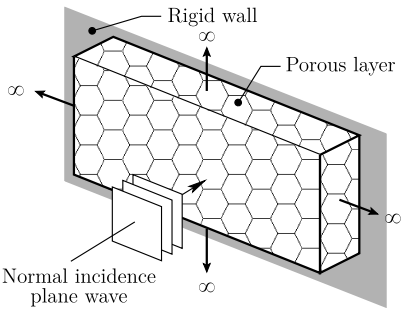

Fig. 5. Poroelastic layer sound absorption test case.

can therefore be employed directly. As the Moore-Neighbour contour tracing algorithm is only valid in a two dimensional context, the routine GetContour() needs to be modified according to the dimensionality of the problem. Computing the contour of an (unstructured) point cloud, i.e. samples in the parameter domain, is a complex problem [49,50] and the sampling effort to detect reasonable patches is likely to grow fast if the parameter dimension is increased. Choosing a parameter sampling method not relying on a Cartesian grid can leverage this additional computational effort to some extent [51,52].

5. Application to bivariate problems

In order to evaluate the potential of the proposed MWCAWE approach, it is applied to two poroacoustic test cases in the following. It is first compared to the Nested Padé Approximants reviewed in Section 3, and then tested in the scope of the multi-patch approach presented in Section 2.2. The implementations in FreeFem++ and Matlab of these examples can be made available upon request. The code and data for the application in Section 5.2 is freely available for download [53].

5.1. Single-patch comparison of the bivariate Nested Padé Approximants and the MWCAWE

The convergence potentials of the bivariate nested Padé Approximants implemented in [4] and the MWCAWE proposed are compared in the following. These are evaluated for the approximation of an academic 1d acoustic absorption problem, already introduced as a test case for nested Padé Approximants in [54]. It consists in a single 2.5 cm-thick poroelastic layer of infinite lateral dimensions, backed by a rigid wall, and is excited by an incident plane wave with normal incidence, see Fig. 5. Assuming a poroelastic material where the motion of the skeleton may be neglected, the Delany-Bazley-Miki empirical model [55,56] is used in order to describe the behaviour of the porous layer as an equivalent fluid. The acoustic response of the porous layer is thus dependent on two independent variables: the frequency f and the static airflow resistivity Φ. The problem, modelled using finite elements, is thus governed by the Helmholtz equation where the density $\widetilde{\rho}_{\mathrm{eq}}(f,\Phi)$ and bulk modulus $\widetilde{K}_{\mathrm{eq}}(f,\Phi)$ are complex-valued and frequency-dependent according to the following algebraic expressions,

$$\widetilde{\rho}_{\mathrm{eq}}(f,\Phi) = \rho_0 \left[1 + \left(5.5 - 8.43i\right) \left(1000 \frac{f}{\Phi} \right)^{-0.632} \right] \left[1 + \left(7.81 - 11.41i\right) \left(1000 \frac{f}{\Phi} \right)^{-0.618} \right], \tag{22a}$$

$$\widetilde{K}_{\mathrm{eq}}(f,\Phi) = \rho_0 c_0^2 \frac{1 + \left(5.5 - 8.43i\right) \left(1000 \frac{f}{\Phi} \right)^{-0.632}}{1 + \left(7.81 - 11.41i\right) \left(1000 \frac{f}{\Phi} \right)^{-0.618}}, \tag{22b}$$

where i is the complex unit, $\tilde{(\cdot)}$ denotes a complex-valued quantity, $\rho_0 = 1.21\ \mathrm{kg\,m^{-3}}$ is the ambient density of the air saturating the pores, and $c_0 = 343\ \mathrm{m\,s^{-1}}$ is the speed of sound in the air.

Proceeding with standard arguments in the FE framework leads to a discretised problem of the form

$$\left[\frac{1}{4\pi^2 f^2 \widetilde{\rho}_{\mathrm{eq}}(f,\Phi)} \mathbf{K}_{\mathrm{p}} - \frac{1}{\widetilde{K}_{\mathrm{eq}}(f,\Phi)} \mathbf{M}_{\mathrm{p}} \right] \mathbf{u} = \mathbf{f}, \tag{23}$$

where $\mathbf{K}_{\mathrm{p}}$ and $\mathbf{M}_{\mathrm{p}}$ are the FE matrices, $\mathbf{u}$ the solution vector containing the pressure fluctuations, and $\mathbf{f}$ the excitation vector. This bivariate problem is thus of the form

$$\mathbf{Z}_{\mathrm{p}}(f,\Phi)\,\mathbf{u}(f,\Phi) = \mathbf{f}, \tag{24}$$

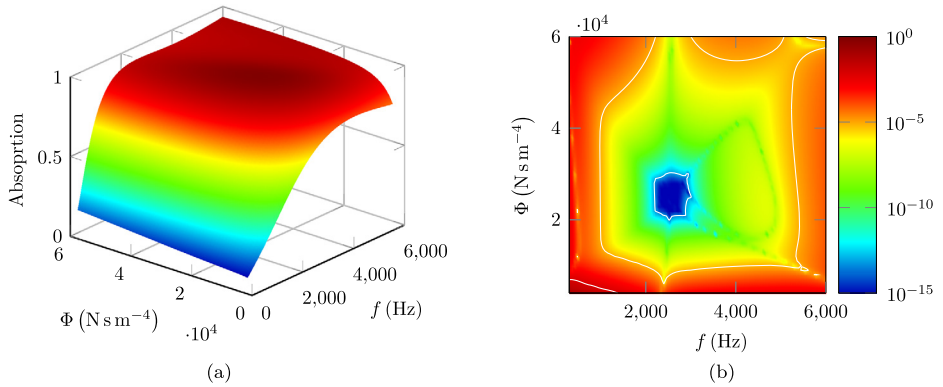

Fig. 6. Nested Padé approximation of the absorption coefficient of a porous material with expansion orders $N_f = N_\Phi = 9$: (a) Surface plot of the absorption coefficient. (b) Relative error in the (f, Φ)-plane view; isocontour of 10^{-13}, 10^{-5}, and 1% relative error plotted in white.

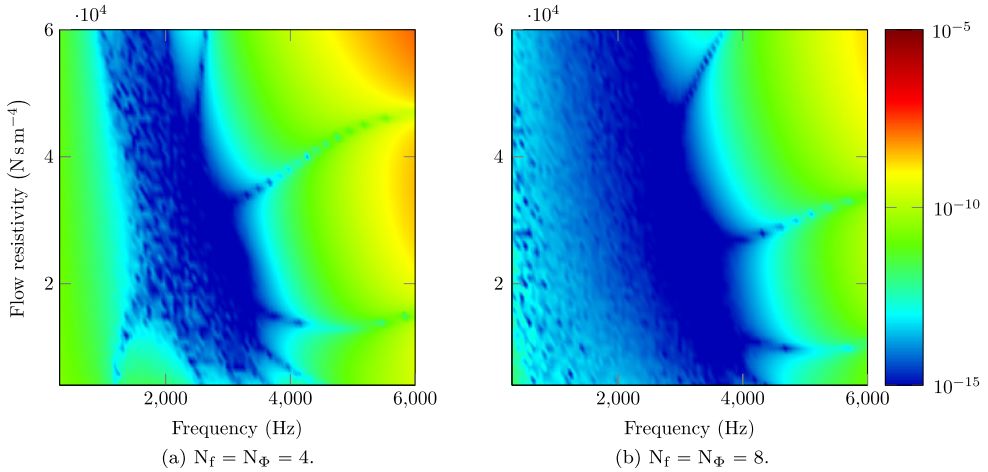

Fig. 7. MWCAWE approximation of the absorption coefficient of a porous material, relative error in the (f, Φ)-plane view.

which corresponds to the generic form of problems introduced in (1). Given that this 1d problem has an analytical solution [54,57], for instance in terms of the absorption coefficient of the poroelastic layer, a comparison of the relative error to this analytical solution is presented in Fig. 6, together with the reference surface response. Nested Padé Approximants, (13), with expansion orders $N_{x_1} = N_f = 9$ and $N_{x_2} = N_\Phi = 9$ for the frequency and flow resistivity, respectively, are considered. Only one patch is evaluated in this case, with a reference point for expansion at $\mathbf{s}_1 = (f_1, \Phi_1) = (2500, 25000)$, a frequency range of $f = [300, 6000]$ Hz and a flow resistivity range of $\Phi = [3000, 60000]$ N s m^{-4}. Fig. 6b highlights a relative error to the analytical solution much lower than 1% (isocontour of highest value among the three plotted: 10^{-13}, 10^{-5}, and 1% relative error) in the almost entire domain of approximation, except at very low values of flow resistivity and low frequencies, where the relative error is more sensitive to small fluctuations due to the low values of the absorption coefficient.

Fig. 7 shows the equivalent to Fig. 6b for the same problem, same expansion point and same domain of approximation, using the MWCAWE approach proposed in Section 2. Although the Nested Padé approach offers the possibility for a component-wise expansion of the solution associated with a single DOF, while projection-based approaches approximate the full solution vector, the improved performance in terms of convergence of the MWCAWE is evident in Fig. 7, as could be expected from, e.g., the comparison of univariate Padé expansions and WCAWE projection-based approach in [18]. Already with expansion orders of $N_f = N_\Phi = 4$ in Fig. 7a, i.e. a truncation at less than half the expansion orders associated with Fig. 6b, the relative error does not exceed 10^{-7} over the entire domain of interest. The accuracy of the reduced-order model is further improved by at least two orders magnitude if the expansion orders of the MWCAWE basis are increased to $N_f = N_\Phi = 8$ in Fig. 7b. This highlights the good convergence properties associated with the MWCAWE approach. Note, that the two-sequence bases associated with the results in Fig. 7 are of dimensions $N_V + (N_V - 2) = (N_f + N_\Phi + 1) + (N_f + N_\Phi - 1) = 2(N_f + N_\Phi)$, taking into account the fact that the second sequence has the same starting and end basis vectors as the former. This is particularly useful in view of the evaluations in the coming sections.

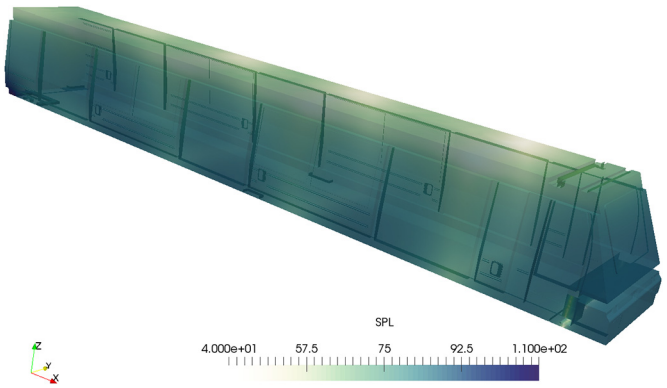

Fig. 8. Sound pressure level of the damped solution at $f = 63$ Hz, $\Phi = 25000\ \mathrm{N\,s\,m^{-4}}$, scale $40 - 110$dB.

5.2. The MWCAWE applied to a large poroacoustic problem

5.2.1. Problem and reference solution

In order to evaluate further the proposed MWCAWE approach and the associated sampling strategy, it is implemented for the approximation of a more complex bivariate poroacoustic problem, adapted from the model introduced in [58]. This problem models the interior cavity of a passenger train equipped with a 15 cm layer of sound absorbing porous material on the top surface. A time-harmonic point source is defined at a lower-back corner of the passenger cavity and parametric sweeps with respect to both the frequency and the static airflow resistivity of the porous layer are performed. All boundary walls are considered as rigid walls, except from the porous boundary. The discretised problem consists of around $\mathrm{N}_n \approx 300\,000$ DOFs. Similarly to the previous example, the porous boundary is modelled by an equivalent fluid formulation, consisting of a modified Helmholtz equation where the equivalent speed of sound $\tilde{c}_\mathrm{p}$ is complex-valued and frequency-dependent, and may be expressed as

$$\tilde{c}_\mathrm{p} = \sqrt{\frac{\widetilde{K}_\mathrm{eq}}{\widetilde{\rho}_\mathrm{eq}}}, \tag{25}$$

adopting the same conventions and notations as for (22).

Proceeding with standard arguments in the FE framework, for two coupled domains governed by the Helmholtz equation (one with a fluid assumed to be inviscid, compressible, homogeneous, and the other with the equivalent fluid model for the fibrous material with high porosity aforementioned), the discretised FE problem may be presented in the following form,

$$\left(\mathbf{K}_\mathrm{a} - \frac{\omega^2}{c_0^2}\mathbf{M}_\mathrm{a} + \mathbf{K}_\mathrm{p} - \frac{\omega^2}{\tilde{c}_\mathrm{p}^{\,2}}\mathbf{M}_\mathrm{p}\right)\mathbf{u} = \mathbf{f}, \tag{26}$$

where $(\cdot)_\mathrm{a}$ denote air cavity global matrices and $(\cdot)_\mathrm{p}$ porous global matrices. $\mathbf{u}$ is the vector of nodal unknowns (acoustic pressure fluctuation here). The right-hand-side vector $\mathbf{f}$, associated with the time-harmonic acoustic excitation is in practice only non-zero at a few DOFs. The FE problem in (26) can again be expressed in the generic form of (1)

$$\mathbf{Z}_\mathrm{pa}(f, \Phi)\,\mathbf{u}(f, \Phi) = \mathbf{f}, \tag{27}$$

and is thus suitable for the proposed multivariate approach.

Fig. 8 illustrates the geometry of the cavity as well as the solution at a given point of the parametric sweep, for a frequency of $f = 63$ Hz and a flow resistivity of $\Phi = 25000\ \mathrm{N\,s\,m^{-4}}$. The reference solution for the bivariate parametric sweep, where $f \in [50, 150]$ Hz and $\Phi \in [3000, 50000]\ \mathrm{N\,s\,m^{-4}}$, at a point in the lower-middle part of a cross-section at the front of the passenger cavity, is plotted in Fig. 9. The solution is evaluated for a uniform distribution of frequency and flow resistivity in these ranges, with increments of 1 Hz and 1000 $\mathrm{N\,s\,m^{-4}}$, respectively, leading to a grid of 4848 sampling points to be evaluated for the domain of interest. The non-damped cavity problem consists of 52 eigenfrequencies in the frequency ranged considered, see [58]. This solution is evidently more complex than the response associated with the 1d absorption problem of Section 5.1, hence a larger reduced model is required to obtain an appropriate accuracy. This reduced model may be sought either in the form of a single point approach with the generation of a large number of basis vectors, or via the multi-point approach proposed in Section 2.2. Both possibilities are considered in the following.

5.2.2. Convergence from a single expansion point

The performance of an MWCAWE basis computed from a single expansion point, i.e. calculating a single local, merged, reduced basis $\mathbf{V}^*_{s_1}$ according to the steps leading to (10), is presented in the following. As mentioned in Section 2.3, as few

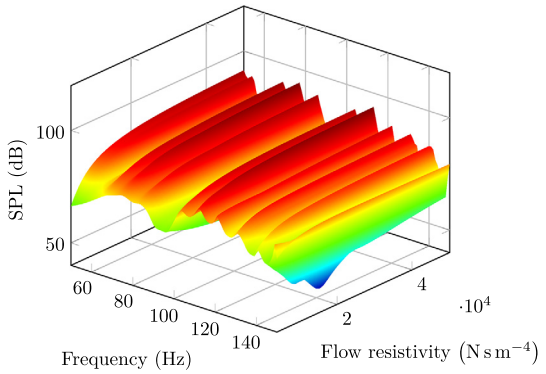

Fig. 9. Reference solution. Sound pressure level at a point in the lower-middle part of a cross-section at the front of the passenger cavity, parametric solution, in dB.

as two sequences are sufficient for a bivariate problem. When introducing a three-sequence approach, only a few of the vectors are discarded by the SVD step in (10) for a threshold $\sigma_{\text{thresh}} = 10^{-15} \cdot \sigma_{\max}$, where $\sigma_{\max}$ corresponds to the largest singular value in the decomposition of (10). In the present case, the two sequences associated with the outer differentiation paths in Fig. 1 are chosen, and the SVD step associated with (9) and (10) subsequently serves the purpose of both merging the bases associated with these sequences and ensuring linearly independent basis vectors.

This two-sequence local basis convergence is tested on an approximation of the reference solution of Fig. 9, with a single approximation patch associated with a reference point arbitrarily chosen at $f = 88$ Hz and $\Phi = 25000\ \mathrm{N\,s\,m^{-4}}$. Given the smoother variations of the sound pressure level (SPL) with respect to flow resistivity when compared to the frequency, it is expected that lower maximum orders of partial derivatives will be required with respect to the flow resistivity than with respect to the frequency. For the sake of illustration of the convergence on this problem, only orders associated with the frequency are increased, while keeping a fixed maximum partial derivative order $\mathrm{N}_\Phi = 8$ for the flow resistivity. Three of these convergence steps are shown in Fig. 10 for maximum partial derivative orders $\mathrm{N}_f = 19$ (Fig. 10, first row), $\mathrm{N}_f = 31$ (Fig. 10, second row), and $\mathrm{N}_f = 55$ (Fig. 10, third row) for the frequency. For each truncation order, the approximate SPL solution over the entire domain is plotted together with the difference (in dB) to the reference solution. The increase from order 19 to order 31 shows a significant improvement of the approximation quality, with a relative difference to the reference solution below 1 dB for more than half the domain of approximation, and only failing to capture the behaviour of the full-order model in the higher end of the spectrum. Increasing this maximum partial derivative order to 55 in frequency brings further improvements in the accuracy, particularly obvious towards the higher frequencies in the spectrum.

As introduced in Section 2.3, alternative derivative pathways, *e.g.* involving univariate sequences (see Fig. 1 and Fig. 2), offer the potential of further reducing the basis sizes for given orders of expansion. Considering the same reference solution, an illustration of the three alternatives introduced in Section 2.3 is plotted in Fig. 11. First Fig. 11a-Fig. 11d allow to compare, for fixed partial derivative orders, the use of bivariate partial derivative sequences (Fig. 1) vs univariate sequences (Fig. 2a) for the same orders. As could be expected, the use of univariate sequences in place of the bivariate alternatives proposed here, without changing the highest orders in each variable, leads to a reduced order of accuracy which echoes the smaller dimension of the projection bases: compare Fig. 11b with its univariate alternative in Fig. 11a, with resulting bases of 20 and 11 vectors, respectively; compare Fig. 11d with its univariate alternative in Fig. 11c, with resulting bases of 36 and 19 vectors, respectively. Of greater interest is the comparison of the approximation resulting from the use of bivariate sequences and univariate alternatives of the type of Fig. 2b, where the overall number of basis vectors is maintained, nevertheless resulting in marginal changes in the approximation: compare Fig. 11d with its univariate alternative in Fig. 11e, with resulting bases of 36 and 37 vectors, respectively. Finally, an additional comparison is presented, where the distribution of the maximum orders of partial derivatives is only redistributed on the most demanding dimension when considering a univariate alternative, where the overall number of basis vectors is maintained. This alternative here leads to improvements in the approximation, increasing the maximum order of partial derivative with respect to frequency from $\mathrm{N}_f = 9$ (Fig. 11d) to $\mathrm{N}_f = 26$ (Fig. 11f), with resulting bases of 36 vectors in both cases. Note however that, despite the potential improvements in accuracy, as observed here, this alternative calls for substantially increased maximum partial derivative orders in one variable, which may not be advantageous in cases where higher order derivatives may need to be symbolically or numerically evaluated, potentially putting a limit on the highest orders available (typically the case here with the expressions of the equivalent density and bulk modulus of the Delany-Bazley-Miki model in (22)).

In view of the single-point convergence tests illustrated here, and in particular due to the fact that the accuracy eventually stagnates with high orders, for which the computational cost associated with the procedure in (5) also grows faster with the number of basis vectors, a multi-point strategy where a trade-off between order of differentiation and number of reference points is required. This trade-off is partly problem-specific, such that in the following, a fixed maximum partial derivative order is chosen to be set a priori, and the multi-point approximation proposed in Section 2.2 is applied to the same train cavity test case.

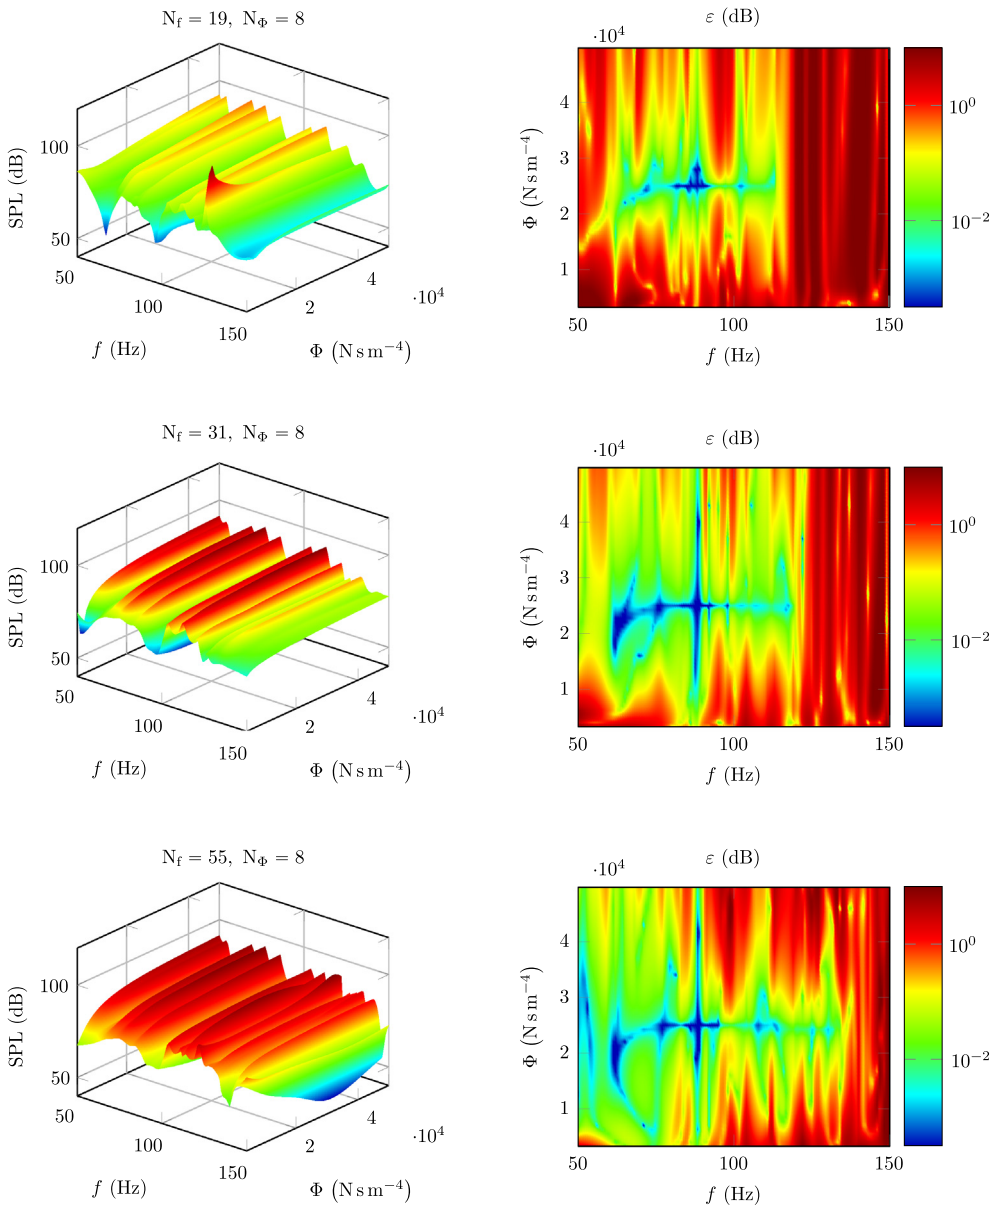

Fig. 10. Convergence, MWCAWE-approximated solution. Reference expansion point: $f = 88$ Hz and $\Phi = 25000\ \mathrm{N\,s\,m^{-4}}$. Different maximum partial derivative orders in frequency/flow resistivity; SPL approximation at the evaluation point (dB, left) and difference to the reference solution (dB, right).

5.2.3. *Multi-point approximation*

In order to illustrate the behaviour of the multi-point procedure introduced in Section 2.2 combined with the residue-based contour-following approach outlined in Section 4.2, fixed orders of maximum partial derivatives are set a priori for all local bases. These are arbitrarily set to orders $N_{x_1} = N_f = 11$ for the frequency and $N_{x_2} = N_\Phi = 5$ for the flow resistivity, thus anticipating the fact that the solution is more sensitive to variations in frequency than flow resistivity in the bounds of interest. As reported for the two-sequence approach, this leads to local merged bases consisting of $2\,(N_f + N_\Phi) = 32$ basis vectors, cf. Sections 2.2 and 2.3 and Fig. 1. The starting reference point is chosen to be the same as for the single-patch convergence test at $f = 88$ Hz and $\Phi = 25000\ \mathrm{N\,s\,m^{-4}}$. For the contour following algorithm, the residual error tolerance is set to $\varepsilon_{\max} = 10^7$, in practice involving a normalisation with respect to the residual error at the reference point for each patch. Finally, the gap tolerance for the iterative introduction of new reference points is set such that the maximum distance between two contours of convergence, $\Delta_{\mathrm{gap\_max}}$, is of 10 units in the sense of the Chebyshev distance, or of 5 units for the distance between a contour and a boundary of the domain. These units correspond, for the considered bivariate problem, to the previously introduced uniform grid associated with frequency and flow resistivity increments. Thus, the domain of interest contains 4848 sampling points. Obviously, the choices of N_{x_1}, N_{x_2}, $\varepsilon_{\max}$ and $\Delta_{\mathrm{gap\_max}}$ are not entirely independent in

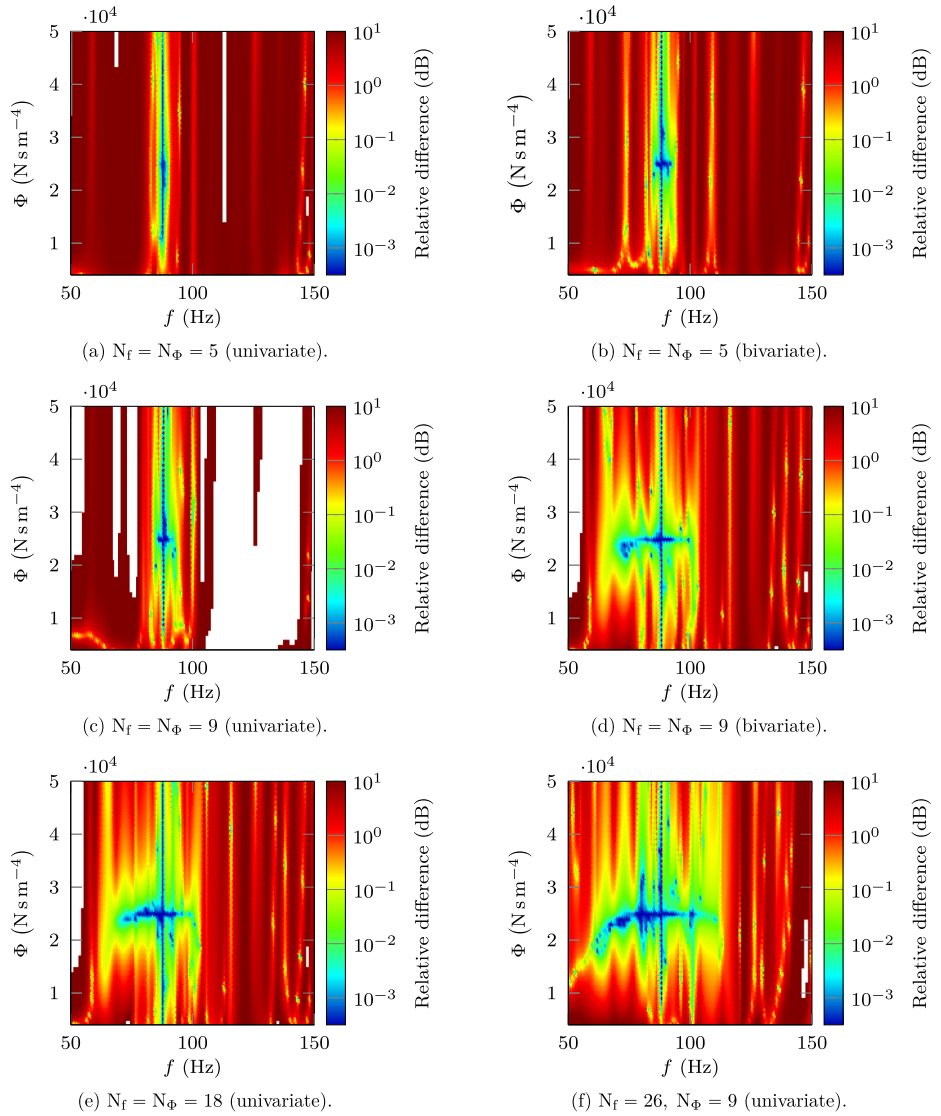

Fig. 11. Comparison of the accuracy with alternative derivation paths, see Fig. 1 and Fig. 2, expansion point: $f = 88$ Hz, $\Phi = 25000\ \mathrm{N\,s\,m^{-4}}$. Difference to the reference solution (dB). Resulting bases of sizes N_{V^*} $(= N_{V^{mer}})$: (a) 11, (b) 20, (c) 19, (d) 36, (e) 37, (f) 36.

relation to the accuracy-efficiency tradeoff, which is a problem-dependent question. However, without any prior knowledge of the problem, it is always possible to set the parameter Δ_{gap_max} to a minimum such that accuracy would be set to highest priority.

Fig. 12 shows four key steps associated with Algorithm 1. The sequence of Fig. 12a to Fig. 12d depicts an increasing number of convergence isocontours associated with local bases independently calculated from the reference points introduced at each iteration of Algorithm 1. The darker contours correspond to the latest addition in the successive iterations of the algorithm. Given the input parameters described above, the sequence ends after nine local bases in Fig. 12d. Note a few gaps remaining between the convergence contours, such as for low values of the flow resistivity in the ranges $[80, 88]$ Hz or $[112, 130]$ Hz. The gaps are in agreement with the introduced gap tolerance Δ_{gap_max}. These nine local bases are however recombined in an SVD step, into a single projection basis, see (18). This recombination is further intended to improve the accuracy of the approximation in the gaps where all the neighbour local bases may contribute. In order both to ensure a well-conditioned reduced system of equations and to reduce the number of DOFs, the basis vectors associated with the lowest singular values are discarded, according to (19), such that $\sigma_{thresh} = 10^{-8} \cdot \sigma_{max}$, where σ_{max} is the largest singular value. In the present case, this truncation results in downsizing the merged basis from $N_{V^{*mer}} = 288$ to $N_{V^{**}} = 275$ basis vectors. This minor reduction of the size of the basis shows the close-to-maximum dimension of the vector space spanned by the independently calculated and merged local bases.

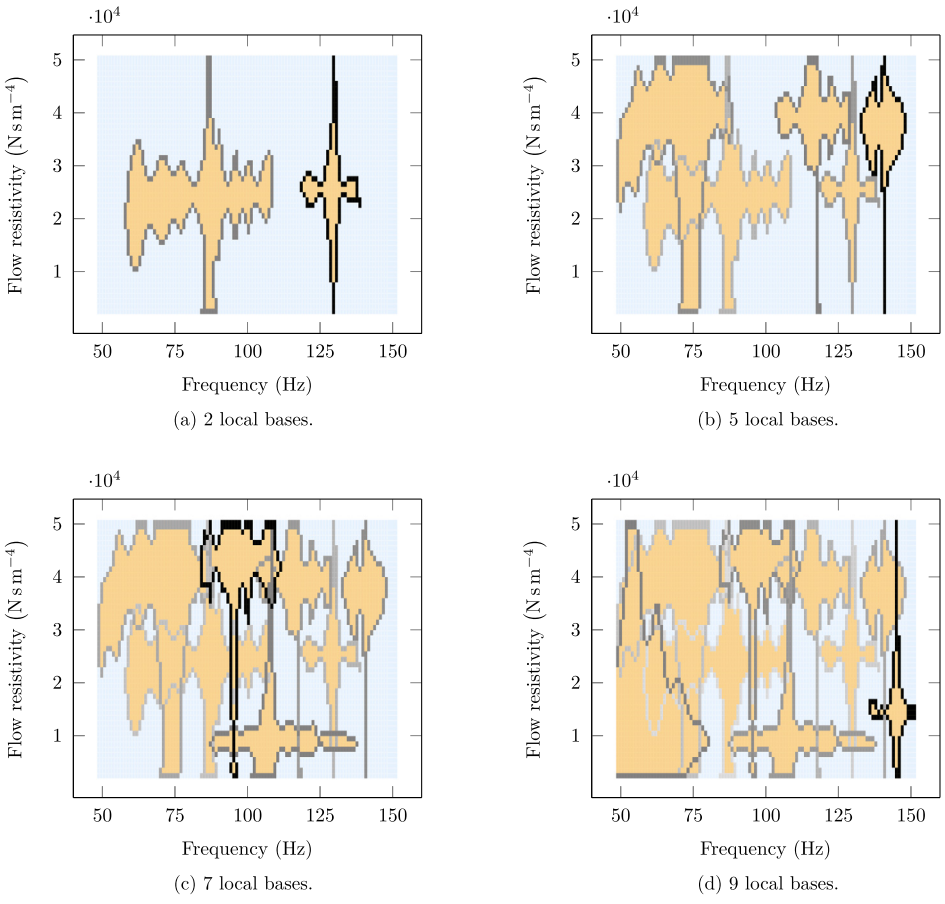

Fig. 12. Contour-following, multi-patch calculation of local bases for the MWCAWE. Starting expansion point: $f = 88$ Hz and $\Phi = 25000$ N s m^{-4}. Fixed partial derivative orders in frequency/flow resistivity: 11/5. Contours of convergence associated with the first 2, 5, 7, and 9 local bases. The regions with darker contours correspond to the latest calculated local bases.

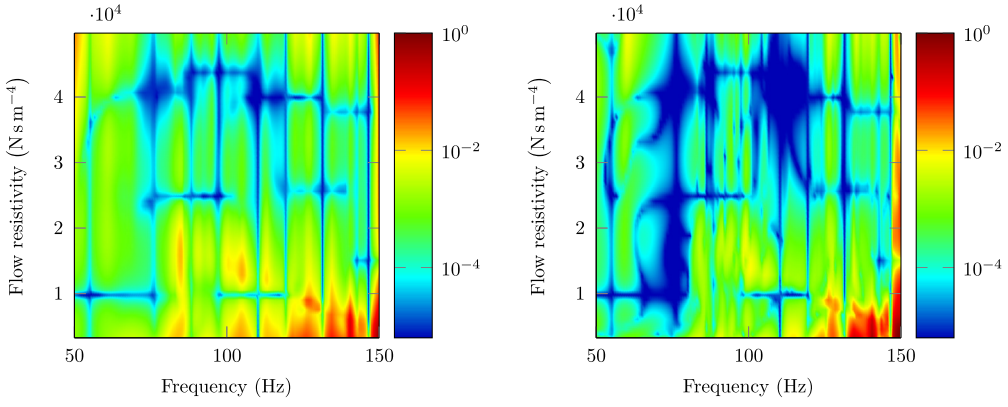

Fig. 13. Error of the approximate solution from the merged bases associated with the contour of Fig. 12. Fixed partial derivative orders in frequency/flow resistivity: 11/5, and 275 final basis vectors from 9 merged local bases.

In Fig. 13 the resulting approximation error associated with the multi-point MWCAWE approach is plotted: the estimated error in Fig. 13a and the actual difference to the reference solution in Fig. 13b. A qualitative comparison between the plots shows the good correlation between the error estimator and the actual approximation error. Thus, the error estimator enables a fast evaluation of the bounds of convergence and the actual resulting approximation error associated with the MWCAWE. It is also noteworthy that, as anticipated in Section 4.3, the error levels in the resulting gaps illustrated in

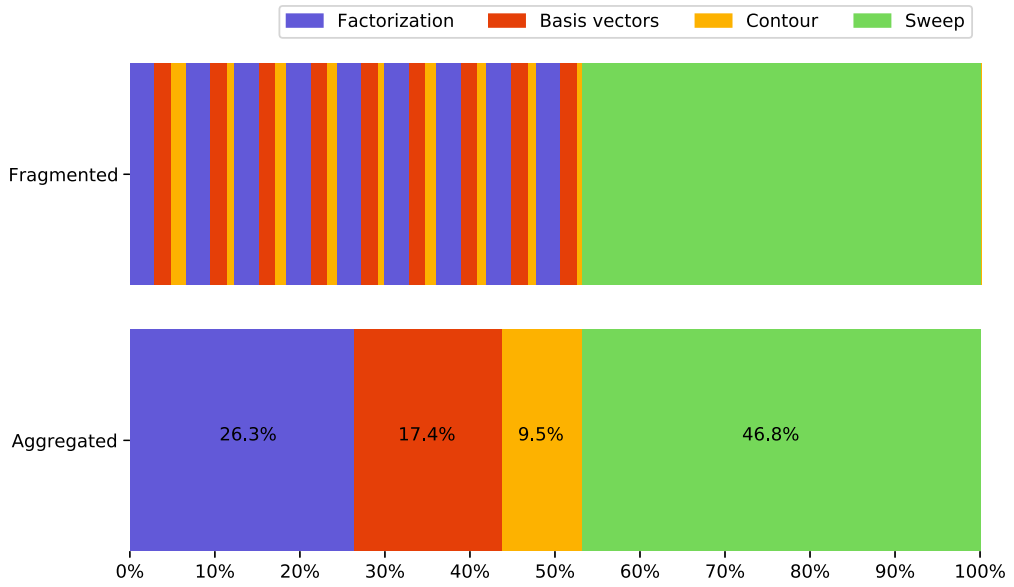

Fig. 14. Evaluation of the distribution of CPU time for the different steps in Algorithm 1.

Fig. 12d, after all nine local bases are taken into consideration and merged into a single global reduced basis, remain very low; for the most parts comparable to the levels within the bounds of convergence determined in Fig. 12. Furthermore, the error levels within these bounds are also several orders lower than the levels reported in Fig. 10, thus highlighting the viability of the multi-point approach in finding a trade-off between size of the resulting reduction basis, computational resources required to establish it, and associated degree of accuracy of the approximation.

In order to illustrate the distribution of the computational time associated with the different steps of the multi-patch approximation, Fig. 14 plots this distribution across four categories:

- *"Factorization"*: the step prior to the start of the MWCAWE recursive basis-vector procedure, (5), where the system matrix at each reference point $\mathbf{s}_n$ is first factorized;
- *"Basis vectors"*: the recursive generation of the vectors of the local bases, as detailed in (5);
- *"Contour"*: the calculation of the boundaries of the patches with the contour-following algorithm, corresponding to the functions GetContour() and Fill() in Algorithm 1;
- *"Sweep"*: the deployment of the recombined set of local bases in order to calculate the low fidelity solution.

The two phases *"Factorization"* and *"Basis vectors"* combined may be associated with the computational time of the function GetBasis() in Algorithm 1. These four phases may be grouped into two groupings: first related to the calculation of the projection bases (*Factorization*, *Basis vectors* and *Contour*), and then the calculation of the low fidelity solution (*Sweep*). The relative computational cost between these phases and groupings may be better observed on the aggregated plot of the estimated distribution of computational cost in Fig. 14. One may conclude that for this specific application, about 10% of the computational time is allocated to the contour retrieving of the patches associated with the nine local bases, while the remaining 90% are about equally distributed between the generation of the nine local bases, and the actual calculation of the low fidelity solution at the 4848 sampling points (*Sweep* phase). Since the *Sweep* phase may be arbitrarily inflated by the actual number of sampling points chosen, a more insightful outcome may be highlighted in the speedup factor per sampling point with respect to the high-fidelity solution, here estimated to be over 600. Note however that this speedup factor is provided solely for reference purposes and may vary significantly based on specific implementation and hardware factors. When considering the grouping of *Factorization*, *Basis vectors* and *Contour*, associated with the preparation of the low-fidelity model, one may highlight for instance that the *Contour* phase corresponds to less than 20% of the overall computational time of this grouping, while almost 50% of that computational time is allocated to the nine factorizations of the system matrix at the reference points. With the fixed orders of expansion chosen here, more resources are thus allocated to the factorization step than to the recursive generation of the basis vectors.

6. Conclusions

The MWCAWE, a multivariate model order reduction method relying on the WCAWE algorithm, is presented. It consists in a two-stage generation of local bases, subsequently recombined for a projection-based approach to approximate the full-order parametric model. At the heart of the method, a reduced set of partial derivation pathways of the parametric system

matrices allows to generate these local bases using the original WCAWE algorithm. The number of partial derivation pathways is in principle linearly increasing with the parametric dimensionality, translating into a quadratic increase in terms of the number of basis vectors. Additionally, in order both to limit the maximum partial derivative orders, associated with a rapid increase of operations associated with the WCAWE algorithm, and to improve the accuracy of the approximation in a given parametric domain, a sampling strategy is proposed. Here presented for the case of bivariate problems, it relies on a residue-based estimation of the approximation error combined with a contour-following algorithm, enabling to sequentially select new points outside the domain of convergence. The MWCAWE approach is shown to have good convergence properties, here compared to a multivariate Padé-based expansion approach on an academic poroacoustic absorption problem. The multi-point strategy is applied to a larger poroacoustic, light train cavity problem, enabling to approximate the entire parametric domain of interest with nine recombined local bases. The associated residue-based estimator is shown to capture the main features of the approximation error. Specific ongoing perspectives to this contribution consist in combining the contour-following approach with an adaptive parameter sampling approach to better leverage the computational effort required to compute patches in the parameter domain where the reduced-order model is considered accurate. This is a prerequisite to make this method applicable to problems with a higher parameter dimension.

CRediT authorship contribution statement

Romain Rumpler: Conceptualization, Data curation, Formal analysis, Funding acquisition, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **Quirin Aumann:** Conceptualization, Data curation, Formal analysis, Investigation, Visualization, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The source code associated with the application in Section 5.2 is available at https://doi.org/10.5281/zenodo.7806380 under the BSD-2-Clause license.

Acknowledgements

The first author gratefully acknowledges the financial support provided by the Swedish Research Council (VR Grants 2015-04925 and 2021-05791), the Centre for ECO^2 Vehicle Design (Sweden's Innovation Agency VINNOVA, Grant Number 2016-05195), and the Swedish Energy Agency (project No. 42100-1).

Appendix A. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jcp.2023.112319.

References

[1] R. Djellouli, C. Farhat, R. Tezaur, A fast method for solving acoustic scattering problems in frequency bands, J. Comput. Phys. 168 (2) (2001) 412–432, https://doi.org/10.1006/jcph.2001.6707.
[2] P. Avery, C. Farhat, G. Reese, Fast frequency sweep computations using a multi-point Padé-based reconstruction method and an efficient iterative solver, Int. J. Numer. Methods Eng. 69 (13) (2007) 2848–2875, https://doi.org/10.1002/nme.1879.
[3] P. Guillaume, Nested multivariate Padé approximants, J. Comput. Appl. Math. 82 (1–2) (1997) 149–158, https://doi.org/10.1016/s0377-0427(97)00081-2.
[4] R. Rumpler, P. Göransson, H.J. Rice, An adaptive strategy for the bivariate solution of finite element problems using multivariate nested Padé approximants, Int. J. Numer. Methods Eng. 100 (9) (2014) 689–710, https://doi.org/10.1002/nme.4777.
[5] P. Benner, S. Gugercin, K. Willcox, A survey of projection-based model reduction methods for parametric dynamical systems, SIAM Rev. 57 (4) (2015) 483–531, https://doi.org/10.1137/130932715.
[6] L. Daniel, O.C. Siong, L.S. Chay, K.H. Lee, J. White, A multiparameter moment-matching model-reduction approach for generating geometrically parameterized interconnect performance models, IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst. 23 (5) (2004) 678–693, https://doi.org/10.1109/tcad.2004.826583.
[7] U. Baur, C. Beattie, P. Benner, S. Gugercin, Interpolatory projection methods for parameterized model reduction, SIAM J. Sci. Comput. 33 (5) (2011) 2489–2518, https://doi.org/10.1137/090776925.
[8] P. Benner, L. Feng, A robust algorithm for parametric model order reduction based on implicit moment matching, in: A. Quarteroni, G. Rozza (Eds.), Reduced Order Methods for Modeling and Computational Reduction, in: Modeling, Simulation and Applications, Springer, Cham, 2014, pp. 159–185.
[9] D. Amsallem, C. Farhat, Interpolation method for adapting reduced-order models and application to aeroelasticity, AIAA J. 46 (7) (2008) 1803–1813, https://doi.org/10.2514/1.35374.
[10] J. Degroote, J. Vierendeels, K. Willcox, Interpolation among reduced-order matrices to obtain parameterized models for design, optimization and probabilistic analysis, Int. J. Numer. Methods Fluids 63 (2) (2010) 207–230, https://doi.org/10.1002/fld.2089.
[11] J. Borggaard, K.R. Pond, L. Zietsman, Parametric reduced order models using adaptive sampling and interpolation, IFAC Proc. Vol. 47 (3) (2014) 7773–7778, https://doi.org/10.3182/20140824-6-ZA-1003.02664.

[12] R. Everson, L. Sirovich, Karhunen–Loève procedure for gappy data, J. Opt. Soc. Am. A 12 (8) (1995) 1657, https://doi.org/10.1364/JOSAA.12.001657.
[13] A.C. Ionita, A.C. Antoulas, Data-driven parametrized model reduction in the Loewner framework, SIAM J. Sci. Comput. 36 (3) (2014) A984–A1007, https://doi.org/10.1137/130914619.
[14] A.C. Rodriguez, L. Balicki, S. Gugercin, The p-AAA algorithm for data driven modeling of parametric dynamical systems, http://arxiv.org/pdf/2003.06536.
[15] R.D. Slone, R. Lee, J.-F. Lee, Well-conditioned asymptotic waveform evaluation for finite elements, IEEE Trans. Antennas Propag. 51 (9) (2003) 2442–2447, https://doi.org/10.1109/TAP.2003.816321.
[16] R.D. Slone, R. Lee, J.-F. Lee, Broadband model order reduction of polynomial matrix equations using single-point well-conditioned asymptotic waveform evaluation: derivations and theory, Int. J. Numer. Methods Eng. 58 (15) (2003) 2325–2342, https://doi.org/10.1002/nme.855.
[17] M.S. Lenzi, S. Lefteriu, H. Beriot, W. Desmet, A fast frequency sweep approach using Padé approximations for solving Helmholtz finite element models, J. Sound Vib. 332 (8) (2013) 1897–1917, https://doi.org/10.1016/j.jsv.2012.05.038.
[18] R. Rumpler, P. Göransson, An assessment of two popular Padé-based approaches for fast frequency sweeps of time-harmonic finite element problems, Proc. Meet. Acoust. 30 (1) (2017) 022003, https://doi.org/10.1121/2.0000649.
[19] X. Xie, H. Zheng, S. Jonckheere, B. Pluymers, W. Desmet, A parametric model order reduction technique for inverse viscoelastic material identification, Comput. Struct. 212 (2019) 188–198, https://doi.org/10.1016/j.compstruc.2018.10.013.
[20] Q. Aumann, G. Müller, Predicting near optimal interpolation points for parametric model order reduction using regression models, PAMM 20 (S1) (2021) e202000352.
[21] T. Bui-Thanh, K. Willcox, O. Ghattas, Model reduction for large-scale systems with high-dimensional parametric input space, SIAM J. Sci. Comput. 30 (6) (2008) 3270–3288, https://doi.org/10.1137/070694855.
[22] S. Chellappa, L. Feng, P. Benner, Adaptive basis construction and improved error estimation for parametric nonlinear dynamical systems, Int. J. Numer. Methods Eng. 121 (23) (2020) 5320–5349, https://doi.org/10.1002/nme.6462.
[23] R. Ullmann, S. Sicklinger, G. Müller, Optimization-based parametric model order reduction for the application to the frequency-domain analysis of complex systems, in: P. Benner, T. Breiten, H. Faßbender, M. Hinze, T. Stykel, R. Zimmermann (Eds.), Model Reduction of Complex Dynamical Systems, in: International Series of Numerical Mathematics, vol. 171, Springer International Publishing, Cham, 2021, pp. 165–189.
[24] C. Jelich, S. Koji Baydoun, M. Voigt, S. Marburg, A greedy reduced basis algorithm for structural acoustic systems with parameter and implicit frequency dependence, Int. J. Numer. Methods Eng. 122 (24) (2021) 7409–7430, https://doi.org/10.1002/nme.6835.
[25] U. Baur, C. Beattie, P. Benner, Mapping parameters across system boundaries: parameterized model reduction with low rank variability in dynamics, PAMM 14 (1) (2014) 19–22, https://doi.org/10.1002/pamm.201410006.
[26] S. van Ophem, E. Deckers, W. Desmet, Parametric model order reduction without a priori sampling for low rank changes in vibro-acoustic systems, Mech. Syst. Signal Process. 130 (2019) 597–609, https://doi.org/10.1016/j.ymssp.2019.05.035.
[27] C. Beattie, S. Gugercin, Z. Tomljanović, Sampling-free model reduction of systems with low-rank parameterization, Adv. Comput. Math. 46 (6) (2020) 83.
[28] C. Beattie, S. Gugercin, Z. Tomljanović, Sampling-free parametric model reduction for structured systems, http://arxiv.org/pdf/1912.11382v1.
[29] B. Peherstorfer, K. Willcox, Dynamic data-driven reduced-order models, Comput. Methods Appl. Mech. Eng. 291 (2015) 21–41, https://doi.org/10.1016/j.cma.2015.03.018.
[30] R. Rumpler, R. Rodríguez Sánchez, P. Göransson, A multivariate, well-conditioned asymptotic waveform evaluation for finite element solutions with complex parametric dependence, in: Proceedings of ISMA 2018 - International Conference on Noise and Vibration Engineering and USD 2018 - International Conference on Uncertainty in Structural Dynamics, 2018, pp. 2403–2414, https://www.scopus.com/inward/record.uri?eid=2-s2.0-85060386599&partnerID=40&md5=617801fa27e35b5d2d93ad401af75d7c.
[31] G.A. Baker Jr., P. Graves-Morris, Padé Approximants, Cambridge University Press, 2007.
[32] C. Beattie, S. Gugercin, Interpolatory projection methods for structure-preserving model reduction, Syst. Control Lett. 58 (3) (2009) 225–232, https://doi.org/10.1016/j.sysconle.2008.10.016.
[33] Z. Bai, Y. Su, SOAR: a second-order Arnoldi method for the solution of the quadratic eigenvalue problem, SIAM J. Matrix Anal. Appl. 26 (3) (2005) 640–659, https://doi.org/10.1137/S0895479803438523.
[34] Y. Saad, Numerical Methods for Large Eigenvalue Problems, Society for Industrial and Applied Mathematics, 2011.
[35] P.J. Bradley, Novel multi-dimensional WCAWE technique for the efficient calculation of RCS, Prog. Electromagn. Res. B 50 (2013) 315–329, https://doi.org/10.2528/pierb13021102.
[36] Q. Li, O. Sigmund, J.S. Jensen, N. Aage, Reduced-order methods for dynamic problems in topology optimization: a comparative study, Comput. Methods Appl. Mech. Eng. 387 (2021) 114149, https://doi.org/10.1016/j.cma.2021.114149.
[37] M. Czarniewska, G. Fotyga, A. Lamecki, M. Mrozowski, Parametrized local reduced-order models with compressed projection basis for fast parameter-dependent finite-element analysis, IEEE Trans. Microw. Theory Tech. 66 (8) (2018) 3656–3667, https://doi.org/10.1109/TMTT.2018.2842744.
[38] D. Panagiotopoulos, E. Deckers, W. Desmet, Krylov subspaces recycling based model order reduction for acoustic BEM systems and an error estimator, Comput. Methods Appl. Mech. Eng. 359 (2020) 112755, https://doi.org/10.1016/j.cma.2019.112755.
[39] V. De La Rubia, U. Razafison, Y. Maday, Reliable fast frequency sweep for microwave devices via the reduced-basis method, IEEE Trans. Microw. Theory Tech. 57 (12) (2009) 2923–2937, https://doi.org/10.1109/TMTT.2009.2034208.
[40] U. Hetmaniuk, R. Tezaur, C. Farhat, An adaptive scheme for a class of interpolatory model reduction methods for frequency response problems, Int. J. Numer. Methods Eng. 93 (10) (2013) 1109–1124, https://doi.org/10.1002/nme.4436.
[41] R. Rumpler, P. Göransson, J.-F. Deü, A finite element approach combining a reduced-order system, Padé approximants, and an adaptive frequency windowing for fast multi-frequency solution of poro-acoustic problems, Int. J. Numer. Methods Eng. 97 (10) (2014) 759–784, https://doi.org/10.1002/nme.4609.
[42] M. Rewieński, A. Lame, M. Mrozowski, et al., A goal-oriented error estimator for reduced basis method modeling of microwave devices, IEEE Microw. Wirel. Compon. Lett. 25 (4) (2015) 208–210, https://doi.org/10.1109/LMWC.2015.2400937.
[43] E. Creixell-Mediante, J.S. Jensen, F. Naets, J. Brunskog, M. Larsen, Adaptive parametric model order reduction technique for optimization of vibro-acoustic models: application to hearing aid design, J. Sound Vib. 424 (2018) 208–223, https://doi.org/10.1016/j.jsv.2018.03.013.
[44] L. Feng, A.C. Antoulas, P. Benner, Some a posteriori error bounds for reduced-order modelling of (non-) parametrized linear systems, ESAIM: Math. Model. Numer. Anal. 51 (6) (2017) 2127–2158, https://doi.org/10.1051/m2an/2017014.
[45] L. Feng, P. Benner, A new error estimator for reduced-order modeling of linear parametric systems, IEEE Trans. Microw. Theory Tech. 67 (12) (2019) 4848–4859, https://doi.org/10.1109/TMTT.2019.2948858.
[46] E.J. Grimme, Krylov Projection Methods for Model Reduction, Dissertation, University of Illinois at Urbana-Champaign, 1997, http://hdl.handle.net/2142/81180.
[47] X. Xie, H. Zheng, S. Jonckheere, A. van de Walle, B. Pluymers, W. Desmet, Adaptive model reduction technique for large-scale dynamical systems with frequency-dependent damping, Comput. Methods Appl. Mech. Eng. 332 (2018) 363–381, https://doi.org/10.1016/j.cma.2017.12.023.
[48] Q. Aumann, G. Müller, Robust error assessment for reduced order vibro-acoustic problems, J. Sound Vib. 545 (2023) 117427, https://doi.org/10.1016/j.jsv.2022.117427.
[49] T. Hackel, J.D. Wegner, K. Schindler, Contour detection in unstructured 3D point clouds, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 1610–1618.

[50] J. Sanchez, F. Denis, D. Coeurjolly, F. Dupont, L. Trassoudaine, P. Checchin, Robust normal vector estimation in 3D point clouds through iterative principal component analysis, ISPRS J. Photogramm. Remote Sens. 163 (2020) 18–35.
[51] S. Chellappa, L. Feng, P. Benner, An adaptive sampling approach for the reduced basis method, in: Realization and Model Reduction of Dynamical Systems, Springer International Publishing, 2022, pp. 137–155.
[52] L. Iapichino, S. Volkwein, Optimization strategy for parameter sampling in the reduced basis method, IFAC-PapersOnLine 48 (1) (2015) 707–712, https://doi.org/10.1016/j.ifacol.2015.05.020.
[53] R. Rumpler, Q. Aumann, Code and data for numerical test case in "MWCAWE: A multivariate WCAWE approach for parametric model order reduction, and a sampling strategy for the bivariate case" (version 1.0), https://doi.org/10.5281/zenodo.7806380, apr 2023.
[54] R. Rumpler, P. Göransson, A finite element solution strategy based on Padé approximants for fast multiple frequency sweeps of multivariate problems, Proc. Meet. Acoust. 19 (1) (2013) 065003, https://doi.org/10.1121/1.4799565.
[55] M. Delany, E. Bazley, Acoustical properties of fibrous absorbent materials, Appl. Acoust. 3 (2) (1970) 105–116, https://doi.org/10.1016/0003-682X(70)90031-9.
[56] Y. Miki, Acoustical properties of porous materials: modifications of Delany-Bazley models, J. Accoust. Soc. Jpn. 11 (1) (1990) 19–24, https://doi.org/10.1250/ast.11.19.
[57] J.-F. Allard, N. Atalla, Propagation of Sound in Porous Media: Modelling Sound Absorbing Materials, 2nd edition, Wiley, Hoboken, NJ, 2009.
[58] R. Rumpler, Padé approximants and the modal connection: towards increased robustness for fast parametric sweeps, Int. J. Numer. Methods Eng. 113 (1) (2018) 65–81, https://doi.org/10.1002/nme.5603.