

A localized reduced-order modeling approach for PDEs with bifurcating solutions[☆]

Martin Hess^a, Alessandro Alla^b, Annalisa Quaini^{c,*}, Gianluigi Rozza^a, Max Gunzburger^d

^a *SISSA, International School for Advanced Studies, Mathematics Area, Trieste, Italy*

^b *PUC-Rio, Department of Mathematics, Rio de Janeiro, Brazil*

^c *Department of Mathematics, University of Houston, Houston, USA*

^d *Department of Scientific Computing, Florida State University, Tallahassee, USA*

Received 22 August 2018; received in revised form 4 March 2019; accepted 25 March 2019

Available online 4 April 2019

Highlights

- Novel ROM approach with localized bases for PDE problems having bifurcating solutions.
- Use of k-means algorithm to cluster snapshots followed by construction of local bases.
- Strategy to find the cluster for a new parameter point and corresponding local basis.
- Effectiveness of the new approach demonstrated by several numerical results.

Abstract

Reduced-order modeling (ROM) commonly refers to the construction, based on a few solutions (referred to as snapshots) of an expensive discretized partial differential equation (PDE), and the subsequent application of low-dimensional discretizations of partial differential equations (PDEs) that can be used to more efficiently treat problems in control and optimization, uncertainty quantification, and other settings that require multiple approximate PDE solutions. Although ROMs have been successfully used in many settings, ROMs built specifically for the efficient treatment of PDEs having solutions that bifurcate as the values of input parameters change have not received much attention. In such cases, the parameter domain can be subdivided into subregions, each of which corresponds to a different branch of solutions. Popular ROM approaches such as proper orthogonal decomposition (POD), results in a global low-dimensional basis that does not respect the often large differences in the PDE solutions corresponding to different subregions. In this work, we develop and test a new ROM approach specifically aimed at bifurcation problems. In the new method, the k-means algorithm is used to cluster snapshots so that within cluster snapshots are similar to each other and are dissimilar to those in other clusters. This is followed by the construction of local POD bases, one for each cluster. The method also can detect which cluster a new parameter point belongs to, after which the local basis

[☆] AA supported by the US Department of Energy Office of Science grant DE-SC0009324. MG supported by the US Air Force Office of Scientific Research grant FA9550-15-1-0001 and US Department of Energy Office of Science grant DE-SC0009324. MH and GR supported by European Union Funding for Research and Innovation through the European Research Council project H2020 ERC CoG 2015 AROMA-CFD project 681447, P.I. Prof. G. Rozza). AQ supported by US National Science Foundation grant DMS-1620384.

* Corresponding author.

E-mail addresses: mhess@sissa.it (M. Hess), alla@mat.puc-rio.br (A. Alla), quaini@math.uh.edu (A. Quaini), gianluigi.rozza@sissa.it (G. Rozza), mgunzburger@fsu.edu (M. Gunzburger).

corresponding to that cluster is used to determine a ROM approximation. Numerical experiments show the effectiveness of the method both for problems for which bifurcation cause continuous and discontinuous changes in the solution of the PDE.

© 2019 Elsevier B.V. All rights reserved.

MSC: 65P30; 35B32; 35Q30; 65N30; 65N35; 65N99

Keywords: Localized reduced bases; Steady bifurcations; Reduced-order modeling; k-means clustering; Proper orthogonal decomposition; Navier–Stokes equations

1. Introduction

Reduced-order modeling has become an indispensable tool for the simulation and control of systems governed by complex, nonlinear partial differential equations (PDEs), as well as for the quantification of uncertainties in outputs of interest that depend on the solutions of those systems. The works cited in this section attest to the broad success such modeling has had in many settings. Here, *reduced-order models* generally refer to inexpensive surrogates for expensive models that are built based on a relatively few solutions of the latter model and for which the expense incurred in the construction process is then amortized over many solutions of the surrogate.

The focus here is on the construction of efficient reduced-order models (ROMs) for PDEs whose solutions bifurcate as the values of parameters appearing in the PDE change. In such cases, using a single ROM results in surrogates that themselves are expensive enough so that their repeated use can incur high costs. Here, we instead build multiple *local ROMs*, each of which can serve to inexpensively determine a set of approximations of PDE solutions that do not span discontinuous or other large changes due to parameter changes. In our methodology, the construction of the multiple local ROMs is automated as is the determination of which local basis one should use for values of the parameters not used in the construction process. A more complete general discussion of ROMs and of our localized ROM approach is given in the rest of this section and details about the latter are provided in subsequent sections.

1.1. Reduced-order modeling in general and for bifurcating solutions

We consider the problem of determining a function $u \in V$ such that

$$\mathcal{N}(u, v; \vec{\theta}) = \mathcal{F}(v; \vec{\theta}) \quad \forall v \in V, \quad \Leftarrow \text{PDE model} \quad (1)$$

where $\vec{\theta} \in \Theta$ denotes a point in a parameter domain $\Theta \subset \mathbb{R}^M$, V a function space, $\mathcal{N}(\cdot, \cdot; \vec{\theta})$ a given form that is linear in v but generally nonlinear in u , and $\mathcal{F}(\cdot)$ a linear functional on V . Note that either or both \mathcal{N} and \mathcal{F} could depend on some or all the components of the parameter vector¹ $\vec{\theta}$. We view the problem (1) as a variational formulation of a nonlinear partial differential equation (PDE) or a system of such equations² in which M parameters appear. However, it is not necessary to do so, e.g., we could treat strong forms of PDEs as well; the notation of (1) allows us to keep the exposition relatively simple.

We are interested in situations in which the solution u of (1) differs in character for parameter vectors $\vec{\theta}$ in different subregions of the parameter domain Θ . Such situations occur if u undergoes bifurcations as $\vec{\theta}$ changes from one subregion to another; in this case we refer to the subdivision of Θ into subregions as a *bifurcation diagram*. We are particularly interested in situations that require solutions of (1) for a set of parameter vectors that span across two or more of the subregions of the bifurcation diagram. Such a situation arises in optimization problems for which an (iterative) optimization algorithm updates the parameter vector at each iteration so that it is possible that the updated parameter vector is in a different subregion than is the previous vector. Uncertainty quantification problems also give rise to the need to find solutions of (1) for possibly many parameter vectors that span across two or more subregions of the bifurcation diagram.

In this paper, we assume that we have a priori knowledge of the bifurcation diagram to allow a more accurate sampling of parameter space around the critical values for the bifurcation. We do so because we focus on

¹ For finite-dimensional spaces, we use the nomenclature “points” and “vectors” interchangeably.

² In Section 3 we use the Navier–Stokes equations as a concrete setting to illustrate our methodology.

the construction and application of a reduced-order model (ROM) for problems with bifurcating solutions. This assumption only affects the first step in the construction of a ROM. In a follow-up paper we treat the case for which no knowledge of the bifurcation diagram is presumed. In Section 4, we discuss what changes are needed so that our approach can handle the more general case.

In general, one cannot solve (1) for u so that one instead chooses an approximating N -dimensional subspace $V_N \subset V$ and then seeks an approximation $u_N \in V_N$ satisfying the discretized system³

$$\mathcal{N}(u_N, v; \vec{\theta}) = \mathcal{F}(v; \vec{\theta}) \quad \forall v \in V_N. \quad \Leftarrow \text{full-order model} \tag{2}$$

Note that if $\{\phi_n(\mathbf{x})\}_{n=1}^N$ denotes a basis for V_N , we have that $u_N = \sum_{n=1}^N u_n \phi_n(\mathbf{x})$ for some set of coefficients $\{u_n\}_{n=1}^N$.⁴

Solving the discretized nonlinear PDE (2) for u_N is often an expensive proposition, especially if multiple solutions are needed. For example, if V_N denotes a finite element subspace constructed with respect to a meshing of nominal grid size h of a d -dimensional spatial domain, we have that $N = O(h^{-d})$ so that the discretized system (2) could be huge. For this reason, one is interested in building surrogates for the solution of (2) that are much less costly to evaluate so that obtaining approximations to the exact solution u of (1) for many choices of the parameter vector $\vec{\theta} \in \Theta$ now becomes feasible. Such surrogates are invariably constructed using a “few” solutions of the expensive, full-order discrete system (2) and can take on many forms such as interpolants and least-squares approximations. Here, we are interested in *reduced-order models* (ROMs) for which one constructs a relatively low-dimensional approximating subspace $V_L \subset V_N$ of dimension L that still contains an acceptably accurate approximation u_L to the solution u of (1). That approximation is determined from the reduced discrete system

$$\mathcal{N}(u_L, v; \vec{\theta}) = \mathcal{F}(v; \vec{\theta}) \quad \forall v \in V_L \quad \Leftarrow \text{global reduced-order model} \tag{3}$$

that, if $L \ll N$, is much cheaper to solve compared to (2). We note that ROM systems are generally dense in the sense that, e.g., matrices associated with (3) are not sparse.

As already stated, the goal is to determine ROM approximations u_L that are acceptably accurate approximations of the solutions u of the continuous model (1). However, because the ROM approximation is constructed through the use of solutions u_N of the spatially discretized problem (2), we have two sources of error in the ROM approximation, an error that can be estimated (in a chosen norm $\|\cdot\|$) by

$$\underbrace{\|u - u_L\|}_{\text{error in ROM approximation}} \leq \underbrace{\|u - u_N\|}_{\text{error due to spatial approximation}} + \underbrace{\|u_N - u_L\|}_{\text{error due to ROM approximation of the spatial approximation}} \tag{4}$$

We mostly focus on the last term in (4). However, one should keep in mind that an efficient overall computational methodology should try to balance the two error terms on the right-hand side of (4) because there is not much sense in having the error due to the ROM be much smaller than the spatial error. This observation may be used to relate L to N so as to provide guidance as to what should be the dimension L of the ROM approximations space.

ROMs in the setting of bifurcating solutions are considered in the early papers [1–4] in the setting of buckling bifurcations in solid mechanics. More recently, in [5] it is shown that a POD approach allows for considerable computational time savings for the analysis of bifurcations in some nonlinear dissipative systems. In [6,7], a reduced basis (RB) method is used to track solution branches from bifurcation points arising in natural convection problems. A RB method is used in [8] to investigate Hopf bifurcations in natural convection problems and in [9] for symmetry breaking bifurcations in contraction–expansion channels. An investigation of symmetry breaking in an expansion channel can be found in [10]. In [11–13], reliable error estimation is used to determine the critical parameter points where bifurcations occur in the Navier–Stokes setting. A recent work on ROMs for bifurcating solutions in structural mechanics is [14]. Finally, we would like to mention that machine learning techniques based on sparse optimization have been applied to detect bifurcating branches of solutions in [15,16] for a two-dimensional laterally heated cavity and Ginzburg–Landau model, respectively.

³ In general, the forms \mathcal{N} and \mathcal{F} themselves are also discretized, e.g., because quadrature rules are used to approximate integrals appearing in their definition. However, here, we ignore such approximations, again to keep the exposition simple.

⁴ Eq. (2) represents a Galerkin type setting in which the trial function u_N and test function v belong to the same space V_N . We could easily generalize our discussion to the Petrov–Galerkin case for which these functions would belong to different approximating spaces.

In most implementations of ROMs, including the papers cited above related to bifurcations, a single *global* basis is used to determine the ROM approximation u_L at any chosen parameter point $\vec{\theta} \in \Theta$ by solving (3). However, in a setting in which Θ consists of subregions for which the corresponding solutions of (1) have different character, such as is the case in the bifurcation setting, it may be the case that L , although small compared to N , may be large enough so that solving the dense ROM system (3) many times becomes a costly proposition. Thus, it seems prudent to construct several *local* bases, each of which is used for parameters belonging to a different subregion of the bifurcation diagram and also possibly to bridge across the boundary between those subregions. Thus, our goal is to construct, say, K such local bases⁵ of dimension L_k , each spanning a local subspace $V_{L_k} \subset V_N$. We then construct K *local reduced-order models*

$$\mathcal{N}(u_{L_k}, v; \vec{\theta}) = \mathcal{F}(v; \vec{\theta}) \quad \forall v \in V_{L_k} \quad \text{for } k = 1, \dots, K \quad (5)$$

\Leftarrow *local reduced-order models*

that provide acceptably accurate approximations u_{L_k} to the solution u of (1) for parameters $\vec{\theta}$ belonging to different disjoint parts of the bifurcation diagram. Of course, if one is to employ several local bases, then one must also determine when one needs to switch from one basis to another.

The use of local ROM bases has been considered in previous works. In [17] snapshots are obtained by sampling the full-order solutions at various time instants; snapshots corresponding to different parameter values are not considered. In [18–20] local bases are determined by projection based clustering instead of Euclidean closeness. In [21,22], k-means clustering and nearest neighbor classifier with respect to parameters or a low-dimensional representation of the current state are used. In [23] k-means clustering is used to generate local bases for use in conjunction with the empirical interpolation method. In [24] local bases determined using k-means clustering and logistic regression classifiers are used for aero-icing problems. In [25] time-based snapshot clustering as well as k-means and a bisection process clusterings of parameter based snapshots are considered and applied to a model problem in cardiac electrophysiology. In [26–28] localized bases are used to resolve fine-scale phenomena in a multiscale setting. In [29] an hp-ROM approach is used to localize ROMs for improved accuracy. In [30] local projection spaces are used in multi-scale turbulence models.

The plan for the rest of the paper is as follows. In Section 2, we discuss in detail the various ingredients that, together, constitute our recipe for the construction and application of a ROM that is well suited for bifurcation problems. Specifically, we discuss how to

1. *select sample points* $\{\vec{\theta}_s\}_{s=1}^S$ *in the parameter domain* Θ ;
2. *compute the corresponding snapshots* $\{u_{N,s}(\mathbf{x})\}_{s=1}^S$, i.e., solutions of the full-order discretized problem (2) for each of the parameter points $\vec{\theta}_s$;
3. *cluster the snapshots so that each cluster corresponds to parameters in a different part of the parameter domain* that could be a subregion of the bifurcation diagram or that could bridge across the boundary of two such subregions;
4. *construct the local bases corresponding to each cluster*;
5. *detect which cluster a new parameter choice belongs to* so that one can use the corresponding local basis to determine a ROM approximate solution.

When these steps are completed, we can solve (5) for the ROM approximation using the appropriate local basis. For the sake of completeness we would like to shortly describe the approach proposed in [15]. The authors deal with time dependent problems with parameter bifurcation. First, they discretize the parameter domain, compute time-dependent snapshots and POD basis for each parameter independently. Then, they collect all the POD basis and select the proper basis, in the online stage, by means of sparse optimization methods. In this way, they are able to detect a subset of the POD basis for the parameter of interest.

In [16], similarly, the authors present a sparse sensing framework to detect bifurcation regimes, through Dynamic Mode Decomposition (DMD, see e.g. [31]) which is a data-driven method that allows to find the best linear fit dynamical systems that approximates the original one. As in [15], they collect all the modes, for any parameter,

⁵ Note that by “local basis” we do not mean local with respect to the support of the basis functions in the spatial domain. In fact, the ROM basis function we construct are global in that respect. By “local basis” we mean a basis that is used to obtain, using (5) ROM solutions for parameters that lie in subregions of the parameter domain.

together and then select the DMD basis by sparse optimization algorithms. The advantage is that with a DMD approach one can learn the dynamics and study eigenvalues and eigenvectors of the linearized system and detect the bifurcation regimes.

In our work, we first collect steady-state snapshots, cluster them and compute POD basis for each cluster. Then, we select the proper basis to use in the online stage through different criteria with respect to [15], as explained in Section 2.5.

In Section 3, to illustrate the implementation and employment of the new algorithm defined in Section 2 in a concrete setting, we consider two problems for the Navier–Stokes equations. Concluding remarks are provided in Section 4.

2. Detailed description of the new localized-basis method

In this section we provide a more detailed description of each of the five steps listed in the recipe given in Section 1.

2.1. Parameter sampling for snapshot generation

The recipe outlined in Section 1 requires, at the start, the selection of S parameter points $\{\vec{\theta}_s\}_{s=1}^S$ in the parameter domain $\Theta \in \mathbb{R}^M$, where M denotes the number of parameters. There exists a large body of literature on methods for this purpose, especially for sampling in hypercubes. Popular sampling techniques include Monte Carlo, quasi-Monte Carlo, Latin hypercube, lattice, orthogonal array, sparse grid, etc. methods. As already mentioned, in this paper we focus on the construction of local reduced bases so that we are content to use crude hand-sampling strategies that pack points near the boundaries between the subregions in the bifurcation diagram. In Section 4, we briefly comment on more sophisticated sampling strategies that potentially result in more accurate ROMs and which we consider in greater detail in our follow-up paper. We remind that the knowledge of the bifurcation diagram allows us to obtain a finer sampling close to the bifurcation points. No a priori knowledge will be assumed at a later stage.

2.2. Generating snapshots

At this point we have in hand a set of S points $\Theta_S = \{\vec{\theta}_s\}_{s=1}^S$ in the parameter domain $\Theta \subset \mathbb{R}^M$. The second task in the recipe given in Section 1 is to determine a solution of the full-order model (2) for each $\vec{\theta}_s$, $s = 1, \dots, S$. The solution corresponding to $\vec{\theta}_s$ is denoted by $u_{N,s}$ and is referred as the *snapshot* corresponding to that parameter point and the set $\mathcal{S} = \{u_{N,s}\}_{s=1}^S$ is referred to as the *snapshot set*.⁶

If $\{\phi_n(\mathbf{x})\}_{n=1}^N$ denotes a basis for V_N , we have, for each $s = 1, \dots, S$, that

$$u_{N,s}(\mathbf{x}) = \sum_{n=1}^N U_{n,s} \phi_n(\mathbf{x}), \quad (6)$$

where the coefficients $\{U_{n,s}\}_{n=1}^N$ are determined by solving the system of N discrete equations (2) for each of the S parameter points $\vec{\theta}_s$, i.e., for each $s = 1, \dots, S$, we solve

$$\mathcal{N}\left(\sum_{n=1}^N U_{n,s} \phi_n(\mathbf{x}), \phi_{n'}; \vec{\theta}_s\right) = \mathcal{F}(\phi_{n'}; \vec{\theta}_s) \quad \text{for } n' = 1, \dots, N. \quad (7)$$

In cases of greatest interest, the PDE problem (1) is nonlinear so that the discrete system (7) is nonlinear as well.

What remains to be specified is the choice of approximating space V_N , the choice of basis $\{\phi_n(\mathbf{x})\}_{n=1}^N$ for V_N , and the choice of solution technique used for the nonlinear discrete system of Eqs. (7). These three choices are usually problem dependent. For the numerical results reported in Sections 3.2 and 3.3, we consider a finite element method [32–34] and a spectral element method [35–37] to effect spatial discretization.

⁶ If the problem (1) is time dependent, then, for each parameter vector sample $\vec{\theta}_s$, one can also include in the snapshot set the solution of (2) evaluated at several chosen instants of time. Here, to present our ideas, it suffices to treat steady-state problems so, of course, we do not include snapshots in time.

Once the approximating space V_N and a basis $\{\phi_n(\mathbf{x})\}_{n=1}^N$ for that space are chosen, a snapshot (6) is completely determined by the vector of its coefficients $\mathbf{U}_{N,s} = (U_{1,s}, \dots, U_{N,s})^T \in \mathbb{R}^N$; naturally, we refer to $\mathbf{U}_{N,s}$ as the *snapshot vector* corresponding to the snapshot function $u_{N,s}(\mathbf{x})$ and the set $\mathcal{U}_{N,S} = \{\mathbf{U}_{N,s}\}_{s=1}^S$ as the *set of snapshot vectors*.

It is important to emphasize that the snapshot set is the only information that is available when constructing the ROM we consider using, as is true for most other ROMs. Thus, if the snapshot set does not contain sufficient information to accurately capture desirable features of the solution of the discrete system (2) for all parameter values, then the ROM will not be able to do so either. As is said colloquially, “if it ain’t the snapshot set, it ain’t in the reduced-order solution”. On the other hand, the snapshot set itself is completely dependent on the choice of sample points $\{\vec{\theta}_s\}_{s=1}^S$.

2.3. Clustering of snapshots

The third step in the recipe given in Section 1 is to cluster the snapshots. It is tempting to believe and to therefore require of any clustering algorithm that any cluster of snapshots should correspond to parameter points belonging to a single subregion in the bifurcation diagram. However, bifurcations come in two guises and only for one guise is such an expectation warranted. In the other, one should not expect or desire that a clustering algorithm performs in this manner.

First, there is the case in which passage of the parameter point from one subregion of the bifurcation diagram to another causes *discontinuous changes* in some characteristic of the solution of (1) or (2), as would happen if one is studying the buckling of a bar under a compressive load; the example in Section 3.3 falls into this category as well. The reasonable expectation in this case is that the snapshots corresponding to parameters in one subregion of the bifurcation diagram are much less “like” the snapshots in other subregions than they are to those in its own subregion. Thus, one would expect the members of any cluster of snapshots to correspond to parameters in a single subregion of the bifurcation diagrams, i.e., that *clustering the snapshots according to their similarity and clustering parameter points according to which subregion of the bifurcation diagram they belong to would yield the same results*. Note that it may happen and it may even be desirable to happen that more than one snapshot cluster corresponds to the same subregion of the bifurcation diagram.

On the other hand, the dependence of the solution on the parameters could be *continuous*, even when passing from one subregion in the bifurcation diagram to another, as would happen if one is studying the stretching of a bar under a tensile load; the example in Section 3.2 falls into this category as well. In this case, the clustering situation is not so clear. One would now in fact reasonably expect that the solution and therefore the snapshots near the boundary of a subregion are more “like” those on the opposite side of that boundary (because the parameter values are close to each other and because of the continuous dependence on the parameters) compared to snapshots in the same subregion but corresponding to parameter values not close to each other. Thus, in this case, a snapshot cluster could correspond to parameter points in more than one subregion in the bifurcation diagram, i.e., *clustering the snapshots according to their similarity and clustering parameter points according to which subregion of the bifurcation diagram they belong to would yield different results*. In the present context, it makes more sense to cluster according to the similarity between snapshots because the ROM bases are built from snapshots so that they influence their construction more directly than the parameters.

2.3.1. *k*-means-based clustering of snapshots

To effect the clustering of the snapshots we use to construct the ROM, we proceed through the following steps, at the heart of which is the use of *k*-means clustering.

- For $k = 2, 3, \dots$, cluster the set of S snapshot vectors $\mathcal{U}_{N,S} = \{\mathbf{U}_{N,s}\}_{s=1}^S$ into k clusters $\{\mathcal{V}_{k'}\}_{k'=1}^k$ using the *k*-means algorithm.
- Determine the corresponding cluster means $\{\mathbf{z}_{k'}\}_{k'=1}^k$.
- For each $k = 2, 3, \dots$, compute the *k*-means variance defined in Section 2.3.2.
- Stop when a plateau is detected in the *k*-means variance.

Note that $k = 1$ is omitted from the iterations because, of course, it corresponds to the entire set of snapshots so that no clustering is needed there is simply only the single mean of all the snapshots.

2.3.2. Determining the optimal number of clusters

A natural question to ask about the clustering of snapshots is how to choose K , the number of clusters? For the same input data, i.e., for the same set of snapshots, the k-means variance

$$V(\mathbf{z}_1, \dots, \mathbf{z}_K; \mathcal{V}_1, \dots, \mathcal{V}_K) = \sum_{k=1}^K \sum_{\{s: U_{N,s} \in \mathcal{V}_k\}} |U_{N,s} - \mathbf{z}_k|^2. \tag{8}$$

decreases as the number of clusters increases. This is obvious because if we have a k-means clustering with K clusters, then adding another cluster mean at one of the data points $U_{N,s}$ would clearly reduce the variance. Furthermore, if $K = S$, then we may pick all the generators to coincide with the data points, leaving us with zero variance but with each cluster containing a single point. Clearly, that is not what one wants, so the question posed above needs a more rational answer.

Variance elbow. If the input data is amenable to clustering, then, for a small number of clusters, the k-means variance reduces quickly as the number of clusters increases. However, eventually the decrease in the variance lessens as one adds more clusters. As an example, consider the discrete data case for which the S input vectors are distributed evenly among four balls of radius 1 but whose centers form a square of side 10. Clearly, going from 2 to 3 k-means clusters would reduce the k-means variance substantially. This is again true when going from 3 to 4 clusters. However, clearly, going from 4 to 5 clusters does not buy one as much reduction in the variance because of the nature of the data.

Thus, if the input data is reasonably clustered, one can expect the plot of the k-means variance vs. the number of clusters to have an “elbow”, i.e., at a certain number of clusters, the plot transitions from having a steep negative slope for a low number of clusters to a gently sloping plateau for a higher number of clusters. Quantitatively, one can use the elbowing effect to choose the number of clusters K to be the smallest integer for which the difference in the k-means variances for $K - 1$ and K clusters is a specified fraction of the difference in the k-means variance between 2 and 3 clusters.

In Sections 3.2 and 3.3, we provide plots of the k-means variance for the examples we consider, plots which exhibit the elbow shape.

2.4. Construction of local ROM bases

The fourth step in the recipe given in Section 1 is to construct local ROM bases, i.e., low-dimensional bases for use in (5) to obtain ROM approximations of the solution of (1), i.e., each local ROM basis is meant to help account for a separate subregion of the parameter domain Θ , including those that straddle across the boundaries between two subregions of the bifurcation diagram.

Perhaps the two approaches in most common use for the construction of ROM bases are the *reduced-basis*⁷ (see, e.g., [38,39]) and *proper orthogonal decomposition* (see, e.g., [40]) methods. RB methods usually use the whole snapshot set as the ROM basis whereas POD bases are determined through an optimization process that selects a basis that spans a subset of the span of the snapshot vectors. A third alternative is to use the cluster means themselves as a ROM basis. Here, so as to have a concrete setting for our numerical illustrations, we choose to use POD bases.

We have in hand the set of snapshot vectors $\mathcal{U}_S = \{U_{N,s}\}_{s=1}^S$ and a clustering $\mathcal{U}_S = \cup_{k=1}^K \mathcal{V}_k$ of that set into disjoint subsets \mathcal{V}_k , $k = 1, \dots, K$, with $|\mathcal{V}_k|$ denoting the cardinality of \mathcal{V}_k ; see Section 2.3. Using this data, we define K local POD bases $\{\psi_{\ell,k}\}_{\ell=1}^{L_k}$ with $L_k \leq |\mathcal{V}_k|$, one for each cluster \mathcal{V}_k , by solving, for $k = 1, \dots, K$, the Euclidean norm minimization problem

$$\min_{\psi_{1,k}, \dots, \psi_{L_k,k}} \sum_{\{s=1: U_{N,s} \in \mathcal{V}_k\}} \left| U_{N,s} - \sum_{\ell=1}^{L_k} (U_{N,s} \cdot \psi_{\ell,k}) \psi_{\ell,k} \right|^2 \tag{9}$$

such that $\psi_{\ell,k} \cdot \psi_{\ell,k'} = \delta_{kk'}$ for $k, k' = 1, \dots, L_K$,

⁷ All ROMs of the type we consider in this paper involve “reduced” bases in the sense that they involve very low-dimensional bases compared to the bases used to effect spatial approximations. However, the nomenclature “reduced basis method” is generally reserved for the context of using all the snapshots as the ROM basis.

where $\delta_{kk'}$ denotes the Kronecker delta. If $K = 1$ there is, of course, no clustering so that the single POD basis is based on all S snapshots; in this case, we refer to the basis as the *global POD basis*.

For each cluster \mathcal{V}_k , $k = 1, \dots, K$, of snapshots, we define the *cluster snapshot matrix* \mathbf{S}_k as the $N \times |\mathcal{V}_k|$ matrix whose columns are the N -dimensional snapshots vectors in cluster \mathcal{V}_k . Then, for each k , the solution of the minimization problem (9) can be obtained from the singular value decomposition (SVD) $\mathbf{S}_k = \mathbf{\Psi}_k \mathbf{\Sigma}_k \mathbf{\Upsilon}_k^T$ of the cluster snapshot matrix. Here, $\mathbf{\Psi}_k$ and $\mathbf{\Upsilon}_k$ are orthogonal $N \times N$ and $|\mathcal{V}_k| \times |\mathcal{V}_k|$ matrices, respectively, and $\mathbf{\Sigma}_k$ is an $N \times |\mathcal{V}_k|$ matrix with all its nonzero entries appearing along the main diagonal in non-increasing order, i.e., $\sigma_{1,k} \geq \sigma_{2,k} \geq \dots \geq \sigma_{|\mathcal{V}_k|,k} \geq 0$, where $\{\sigma_{\ell,k}\}_{\ell=1}^{|\mathcal{V}_k|}$ denote the singular values of \mathbf{S}_k . The columns of $\mathbf{\Psi}_k$ and $\mathbf{\Upsilon}_k$ are referred to as the *left and right singular vectors* of \mathbf{S}_k , respectively, and the nonzero entries of $\mathbf{\Sigma}_k$ are referred to as the *singular values* of \mathbf{S}_k . For each k , the solution of the minimization problem (9) is given by the first L_k columns of $\mathbf{\Psi}_k$, i.e., *the L_k local POD basis vectors corresponding to the cluster \mathcal{V}_k are given by $\{\psi_{\ell,k}\}_{\ell=1}^{L_k}$, where $\psi_{\ell,k}$ is the ℓ -th left-singular vector of \mathbf{S}_k .*

Of course, each POD basis vector induces a function in the space V_N used for the spatial discretization (2) of the given problem (1). Again letting $\{\phi_n(\mathbf{x})\}_{n=1}^N$ denote a basis for V_N , we have, corresponding to the cluster \mathcal{V}_k , the POD basis functions

$$v_{\ell,k}(\mathbf{x}) = \sum_{n=1}^N \psi_{n,\ell,k} \phi_n(\mathbf{x}) \quad \text{for } \ell = 1, \dots, L_k, \tag{10}$$

where $\psi_{n,\ell,k}$ denotes the n th component of the POD basis vector $\psi_{\ell,k}$. Then, the subspace of the spatial approximation space V_N corresponding to the cluster \mathcal{V}_k in snapshot space is given by $V_{L_k} = \text{span}\{v_{\ell,k}\}_{\ell=1}^{L_k}$. Note that the error in (9) is given by the sum of the neglected singular values $\sigma_{\ell,k}$ for $\ell = L_k + 1, \dots, |\mathcal{V}_k|$. Hence, one might choose the number of basis functions L_k such that $\sum_{\ell=L_k+1}^{|\mathcal{V}_k|} \sigma_{\ell,k}$ is smaller than some prescribed tolerance. We further note that this is a heuristic consideration because the problem (9) is dependent on the snapshots.

2.5. Determining which local basis to use for a new parameter

In the fifth step in the recipe given in Section 1, the task faced is to identify, for any parameter point $\vec{\theta} \in \Theta$ which was not among those used to generate the snapshots, which local basis should be used in (5) to determine the corresponding ROM approximation. The local bases are associated with the clustering of snapshots so seemingly, to identify which local basis one should use, one should solve the expensive full-order model (2) for that $\vec{\theta}$ and then determine which of the K cluster means of the snapshot clusters $\{\mathcal{V}_k\}_{k=1}^K$ that solution is closest to. Obviously one does not want to do this because involves costs depending on the number of full-order degrees of freedom. Thus, *although the snapshots are used to determine the clusters, the identification of the local basis to be used has to be done in parameter space.*

Thus the first step in any recipe for assigning new parameters to clusters is to *determine the parameter clustering induced by the snapshot clustering*. There is, of course, a one-to-one correspondence between a snapshot vector $\mathbf{U}_{N,s}$ (or the corresponding snapshot function $u_{N,s}(\mathbf{x})$) and the parameter point $\vec{\theta}_s$ used to determine that snapshot. Thus, a clustering $\{\mathcal{V}_k\}_{k=1}^K$ of the snapshots trivially induces a clustering of the sample points $\Theta_S = \{\vec{\theta}_s\}_{s=1}^S$. Specifically, we have that $\Theta_S = \cup_{k=1}^K \Theta_k$, where, for $k = 1, \dots, K$,

$$\Theta_k = \{\vec{\theta}_s \in \Theta_S : \mathbf{U}_{N,s} \in \mathcal{V}_k\}. \tag{11}$$

There are potential difficulties stemming from the fact that the clustering is done with respect to the snapshots whereas the identification of which basis to use is done with respect to the parameters. For example, a tempting way to assign a new parameter to a cluster (and thus select which local basis to use for solving the ROM model for that parameter) is to first determine the cluster means in parameter space and then assign the new parameter to the cluster corresponding to the closest cluster mean. However, a parameter cluster mean generally does not correspond to the snapshot cluster mean, i.e., if one were to solve the full-order model using the parameter cluster mean, the resulting full-order solution would be different from the snapshot cluster mean. As a result, the possibility exists that the wrong local ROM basis would be assigned to a new parameter. This is much more likely to occur in the case where transitions across boundaries in the bifurcation diagram result in continuous solutions and it is likely that clusters span across such boundaries. In this case, because the transition results in continuous solutions, the ROM approximation error when using the “wrong” local basis still be tolerable, i.e., less than 1%. In the rarer case

in which the transition across boundaries in the bifurcation diagram is discontinuous, using the “wrong” local basis can be expected to result in relatively large approximation error.

In one dimension, i.e., for the case of a single parameter, an alternate means for assigning new parameters to parameter clusters makes use of the cluster midrange and cluster radius and proceeds as follows. We have in hand the clustering $\{\mathcal{V}_k\}_{k=1}^K$ of the snapshots determined as described in Section 2.3 and a point $\bar{\theta} \in \Theta$ for which we want to obtain the corresponding ROM solution. Then, we

- determine, using (11), the clustering $\{\Theta_k\}_{k=1}^K$ in the parameter domain induced by the clustering of snapshots;
- compute, for $k = 1, \dots, K$, the midrange

$$M_k^{(mr)} = \frac{1}{2} \left(\min_{\{s=1, \dots, S, \theta_s \in \Theta_k\}} \theta_s + \max_{\{s=1, \dots, S, \theta_s \in \Theta_k\}} \theta_s \right)$$

and the radius

$$\bar{r}_k = M_k^{(mr)} - \min_{\{s=1, \dots, S, \theta_s \in \Theta_k\}} \theta_s$$

of the parameters in each parameter cluster;

- assign the new parameter to the parameter cluster Θ_k such that

$$|\theta - M_k^{(mr)}| - \bar{r}_k < |\theta - M_{k'}^{(mr)}| - \bar{r}_{k'} \quad \text{for } k' = 1, \dots, K, \quad k' \neq k,$$

i.e., the parameter cluster whose midrange minus the cluster radius is closest to that parameter, note that the term $|\bar{\theta} - \bar{m}_k| - \bar{r}_k$ can also be negative;

- the parameter cluster Θ_k so identified corresponds to a specific snapshot cluster \mathcal{V}_k ;
- the local basis to use is $\{\psi_{\ell,k}\}_{\ell=1}^{L_k}$, i.e., the local basis corresponding to snapshot cluster \mathcal{V}_k .

Limited computational experiments indicate that this may be a superior avenue for assigning a local basis to a new parameter when compared to using parameter cluster means. Of course, there is no shortage of ways to generalize the concepts of midrange and radius to higher dimensions, i.e., to the case of multiple parameters. For example, for the midrange, one could choose the center of the smallest ball containing all the parameters within a cluster, or one could average the positions of the vertices of the convex hull polytope of the parameters, or one could determine the midrange for each component of the parameter vector and use the vector of midranges as the testing point. The corresponding radii would be the radius of the ball, the distance from the midrange to the furthest vertex, and the largest of the one-dimensional radii. In one dimension, all three produce the midrange point.

3. Application to the incompressible Navier–Stokes equations

We use a variational formulation of the incompressible Navier–Stokes equations as the concrete setting of (1) to illustrate our methodology. We first define the Navier–Stokes problem and then consider examples for both continuous and discontinuous transitions through bifurcation points. For the first case, we consider two types of spatial discretizations. Together, the examples are meant to illustrate the robustness of our approach with respect to both solution behaviors and spatial discretization choices.

3.1. Problem definition

The Navier–Stokes equations for incompressible flows are given by

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla \mathbf{u}) - \nu \Delta \mathbf{u} + \nabla p &= \mathbf{f} & \text{in } \Omega \times (0, T] \\ \nabla \cdot \mathbf{u} &= 0 & \text{in } \Omega \times (0, T], \end{aligned} \tag{12}$$

where \mathbf{u} and p denote the unknown velocity and pressure fields, respectively, $\nu > 0$ denotes the kinematic viscosity of the fluid, and \mathbf{f} accounts for possible body forces. These equations describe the incompressible motion of a viscous, Newtonian fluid in the spatial domain $\Omega \subset \mathbb{R}^d$, $d = 2$ or 3 , over a time interval of interest $(0, T]$. The boundary of Ω is denoted by $\partial\Omega$. If the fluid acceleration becomes negligible, i.e., the system has evolved towards a steady state, the time-derivative term in (12) can be omitted.

To complete the definition of the flow problem, (12) needs to be endowed with initial and boundary conditions, which, for the examples of Sections 3.2 and 3.3, take the form

$$\mathbf{u} = \mathbf{u}_0 \quad \text{in } \Omega \times \{0\} \quad (13)$$

$$\mathbf{u} = \mathbf{u}_D \quad \text{on } \partial\Omega_D \times (0, T] \quad (14)$$

$$-p\mathbf{n} + \nu \frac{\partial \mathbf{u}}{\partial \mathbf{n}} = \mathbf{g} \quad \text{on } \partial\Omega_N \times (0, T], \quad (15)$$

where $\partial\Omega_D \cap \partial\Omega_N = \emptyset$ and $\overline{\partial\Omega_D} \cup \overline{\partial\Omega_N} = \overline{\partial\Omega}$. Here, \mathbf{u}_0 , \mathbf{u}_D , and \mathbf{g} are given and \mathbf{n} denotes the unit normal vector on the boundary $\partial\Omega_N$ directed outwards. In the rest of this section, we will explicitly denote the dependence of the solution of the problem (12)–(15), and possibly of \mathbf{f} , on the parameter vector $\vec{\theta}$.

To define a variational formulation of system (12)–(15), we let $L^2(\Omega)$ denote the space of square integrable functions in Ω and $H^1(\Omega)$ the space of functions belonging to $L^2(\Omega)$ with first derivatives in $L^2(\Omega)$. We then define the sets

$$\mathbf{V} := \{\mathbf{v} \in [H^1(\Omega)]^d : \mathbf{v} = \mathbf{u}_D \text{ on } \partial\Omega_D\}, \quad \mathbf{V}_0 := \{\mathbf{v} \in [H^1(\Omega)]^d : \mathbf{v} = \mathbf{0} \text{ on } \partial\Omega_D\}.$$

Then, the standard variational form corresponding to (12)–(15) is given as follows: find $(\mathbf{u}(\vec{\theta}), p(\vec{\theta})) \in \mathbf{V} \times L^2(\Omega)$ satisfying the initial condition (13) and

$$\begin{aligned} \int_{\Omega} \frac{\partial \mathbf{u}(\vec{\theta})}{\partial t} \cdot \mathbf{v} \, d\mathbf{x} + \int_{\Omega} (\mathbf{u}(\vec{\theta}) \cdot \nabla \mathbf{u}) \cdot \mathbf{v} \, d\mathbf{x} - \int_{\Omega} p(\vec{\theta}) \nabla \cdot \mathbf{v} \, d\mathbf{x} \\ = \int_{\Omega} \mathbf{f}(\vec{\theta}) \cdot \mathbf{v} \, d\mathbf{x} + \int_{\partial\Omega_N} \mathbf{g} \cdot \mathbf{v} \, d\mathbf{x} \quad \forall \mathbf{v} \in \mathbf{V}_0 \\ \int_{\Omega} q \nabla \cdot \mathbf{u}(\vec{\theta}) \, d\mathbf{x} = 0 \quad \forall q \in L^2(\Omega). \end{aligned} \quad (16)$$

Problem (16) constitutes the particular case of the abstract problem (1) we use for the numerical illustrations.

The nonlinearity in problem (16) can produce a loss of uniqueness of the solution, with multiple solutions branching from a known solution at a bifurcation point. We consider, in Sections 3.2 and 3.3, two specific problems for which the parameter domain Θ contains two or more regions across whose common boundaries steady bifurcations occur. Because we are interested in studying flow problems close to steady bifurcation points, our snapshot sets include only steady-state solutions as explained in Section 2.2. To obtain the snapshots, we approximate the solution of problem (16) by a time-marching scheme that we stop when sufficiently close to the steady state, e.g., when the stopping condition $\frac{\|\mathbf{u}_N^n - \mathbf{u}_N^{n-1}\|_{L^2(\Omega)}}{\|\mathbf{u}_N^n\|_{L^2(\Omega)}} < \text{tol}$ is satisfied for a prescribed tolerance $\text{tol} > 0$, where n denotes the time-step index.

3.2. Example 1: Jet emanating from an orifice into a channel

We first consider a classical benchmark test: a jet exiting through an orifice into a rectangular channel. The simplicity of the geometry that is partially characterized by the expansion ratio, i.e., the ratio of the orifice size to the channel height, but which still results in complex flow have made this problem a popular choice for use in testing computational models; see, e.g., [41–44] and the references cited therein.

In (12), we set $\mathbf{f} = \mathbf{0}$. The channel height is set to 6. At the orifice, which is of size $w = 1$ and is centered on the left side of the channel, we impose (14) with \mathbf{u}_D having a parabolic horizontal velocity component with maximum horizontal speed $U = 1/4$ and zero vertical component. At the remainder of the left end of the channel and at the top and bottom of the channel, we apply (14) with $\mathbf{u}_D = \mathbf{0}$. The right end of the channel is an outlet at which (15) is imposed with $\mathbf{g} = \mathbf{0}$. With this condition along with the other data imposed, if the channel is sufficiently long, the outflow will be fully established [45]. For this reason, we choose the channel length to be 36. The Reynolds number $Re = Uw/\nu$ can be used to characterize the flow regime; it can be thought of as the ratio of inertial forces to viscous forces. For large values of the Reynolds numbers, inertial forces are dominant over viscous forces and vice versa. For the data chosen above, we have that $Re = 1/(4\nu)$.

For a fixed expansion ratio, as the Reynolds number Re increases from zero, the sequence of flow configurations is as follows. For sufficiently small values of Re , a steady *symmetric* jet is observed. Moffatt eddies form close to




Fig. 1. The 32 spectral elements used for spatial approximation which results in 14,259 degrees of freedom.

the corners downstream of the expansion. As the inertial effects of fluid become more important, the Moffatt eddies develop into two recirculation regions of equal size. As Re increases further, flow symmetry about the horizontal centerline is initially maintained and the recirculation length progressively increases. However, at the critical value $Re \approx 33$, one recirculation zone expands whereas the other shrinks, giving rise to a steady *asymmetric* jet. This asymmetric solution remains stable as Re increases further, but the asymmetry becomes more pronounced, as shown in [44]. The configuration with a symmetric jet is still a solution, but is unstable [46]. Snapshots of the flow fields for three stable scenarios are illustrated in Section 3.2.2. Changes in the solution as Re increases are *continuous*, even when passing from a symmetric to an asymmetric configuration. This loss of symmetry in the steady solution as Re changes is a supercritical pitchfork bifurcation [47].

3.2.1. Spatial discretization approaches

For this example, we compare the use of two spatial discretization methods, specifically, a finite element method (FEM) and a spectral element method (SEM). Our goal here is to show the possible effect that spatial discretization choices can have on the results obtained for our ROM approach.

For the *spectral element discretization*, the SEM software framework Nektar++, version 4.4.0, (see <https://www.nektar.info/>) is used for the full-order simulations. The domain is discretized into 32 quadrilateral elements as shown in Fig. 1. Modal Legendre ansatz functions of order 12 are used in every element and for every solution component. This results in 4753 global degrees of freedom for each of the horizontal and vertical velocity components and the pressure for the time-dependent simulations. For temporal discretization, an IMEX scheme of order 2 is used with a time-step size of $\Delta t = 10^{-3}$; typically 10^6 time steps are needed to reach a steady state.

For the *finite element discretization*, the FEM software framework FEniCS (see <https://fenicsproject.org/>) is used for the full-order simulations. We effect spatial discretization using the Taylor–Hood finite element pair, i.e., continuous piecewise quadratic polynomials for approximating the velocity components and continuous piecewise linear polynomials for approximating the pressure, all based on the same grid. The triangular grid used results in 90,876 degrees of freedom. For the temporal discretization, a modified version of Chorin’s method, referred to as the incremental pressure correction scheme (IPCS) [48], is used with a time step $\Delta t = 0.1$.

The numerical study shows that a supercritical pitchfork bifurcation occurs around the critical Reynolds number ≈ 33 . For values of Re lower than that critical value, a single branch of symmetric stable approximate solutions are obtained. For values above the critical value, three different steady state solutions are possible. Besides a solution which is symmetric with respect to the horizontal midline of the domain, two other branches of approximate solutions that are asymmetric with respect to the horizontal midline can occur. The symmetric approximate solution is now unstable whereas the asymmetric approximate solutions are stable. This is all in accordance to what is known about the exact solution.

To obtain the unstable symmetric approximate solutions for higher Reynolds numbers, a continuation method is used in which the solution is determined with an initial condition given from a solution at a slightly lower Reynolds number. To transition onto one of the two stable branches, a solution is first computed with a non-symmetric inflow condition, i.e., a stronger inflow towards the either the upper or lower wall, as desired. This solution is then used as an initial condition with a symmetric inflow condition, as specified in the model description, to obtain an approximate solution on the desired branch. The desired branch is then followed by again using a continuation method to determine approximate solutions for higher values of the Reynolds numbers. Using these strategies, the full pitchfork bifurcation diagram can be obtained as illustrated in Fig. 2 that provides that diagram based on the point value of the approximate vertical velocity component at the point (1.0, 3.0).

When computing the snapshots needed to build a reduced-order model, one clearly needs to follow one or the other of the asymmetric branches i.e., to not mix snapshots obtained from both branches. In our calculations, we use the lower branch of asymmetric solutions as determined using the continuation strategy mentioned above.




Fig. 2. For the jet problem, the full bifurcation diagram, based on the vertical velocity component at the point (1.0, 3.0) for Reynolds numbers from 5 to 100.




Fig. 3. Top to bottom: full-order steady-state streamlines for $Re = 10$, $Re = 35$, and $Re = 100$ given by FEM (left) and SEM (right). Colors correspond to the magnitude of the velocity according to the given scales. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

3.2.2. Results for the offline stage: the construction of local ROM bases

We next compare results obtained using the two spatial discretization methods for the offline stage in which local ROM bases are constructed.

Snapshots. We computed 57 snapshots associated with the following Reynolds numbers:

$$\begin{aligned}
 Re = \{ & 5, 5.71, 6.67, 8, 10, 11.77, 14.29, 18.18, 20.83, 22.73, 25, 26.32, 27.78, 29.41, \\
 & 31.25, 32.26, 32.79, 33.33, 33.9, 34.4828, 35.0877, 35.7143, 36.36, 37.03, 37.73, \\
 & 38.46, 39.22, 40, 40.82, 41.67, 42.55, 43.48, 44.44, 45.45, 46.51, 47.62, 48.78, \\
 & 50, 52.63, 55.56, 58.82, 62.5, 66.67, 68.96, 71.43, 74.07, 76.93, 80, 83.33, 84.75, \\
 & 86.95, 89.28, 90.91, 92.59, 95.24, 98.04, 100 \}.
 \end{aligned} \tag{17}$$

We recall that the snapshots are the solution of (12) close to steady state for each Re . Note that the sampling is finer around the critical value for the bifurcation reported in the literature (i.e., around $Re = 33$).

First, in Fig. 3, we provide visual comparisons of steady-state streamline snapshots obtained using the two methods and for three values of the Reynolds number: $Re = 10$ which is below the critical Reynolds number and $Re = 35$ and 100 which are slightly above and well above, respectively, the critical Reynolds number. Not surprisingly, the results for two methods are in very good agreement, with the visible differences occurring in very low velocity regions which is unavoidable because the discretization errors of the schemes are most probably larger than the speed in those regions.

Bifurcation diagram. We next compare, in Fig. 4, the (lower half of the) bifurcation diagram for the vertical velocity u_y at the point (1, 3) as obtained by the SEM and FEM methods. The two diagrams are in very close agreement, with the only visual differences occurring for Reynolds numbers such that the flow is only very slightly asymmetric




Fig. 4. Lower stable branch of the bifurcation diagram over Reynolds numbers from 5 to 100 for the jet problem for the FEM (red) and SEM (blue) discretizations. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)




Fig. 5. k-means variance versus number of clusters for the jet problem for the SEM (blue crosses) and FEM (red crosses) discretizations. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

and thus the vertical velocity is extremely small. Again, this discrepancy is unavoidable because the discretization errors of the schemes are most probably larger than $|u_y(1, 3)|$.

k-means variance. The k-means variance is plotted vs. the number of clusters in Fig. 5 for both spatial discretization approaches. We use the elbowing effect of the k-means variance to determine the optimal number of clusters so the close agreement between the SEM and FEM results would result in the same choice for the optimal number of clusters. With regard to the plot itself, it is clear that the value of the k-means variance reduces quickly for a very few clusters and that the reduction lessens as the number of clusters increases. The decision about how many clusters to use should be guided by the goal of making the clustering error commensurate with the other errors incurred in the ROM process.

Clustering. We apply the k-means method to cluster the snapshots produced by the two spatial discretization approaches. Results are shown in Fig. 6 for 3, 4, and 5 clusters. The k-means method decides on how to cluster the snapshots by minimizing the total clustering variance. Fig. 5 hints that two discretization methods will result in very similar clusterings; this is confirmed in Fig. 6 for which one observes very close agreement between the clusterings obtained using the two methods. The snapshots for 3, 4, 5 clusters are divided as follows:

$$\begin{aligned}
 Re &= \{Re_i\}_{i=1}^{25} \cup \{Re_i\}_{i=26}^{35} \cup \{Re_i\}_{i=36}^{57}, \\
 Re &= \{Re_i\}_{i=1}^9 \cup \{Re_i\}_{i=10}^{26} \cup \{Re_i\}_{i=27}^{42} \cup \{Re_i\}_{i=43}^{57}, \\
 Re &= \{Re_i\}_{i=1}^9 \cup \{Re_i\}_{i=10}^{25} \cup \{Re_i\}_{i=26}^{39} \cup \{Re_i\}_{i=40}^{47} \cup \{Re_i\}_{i=48}^{57},
 \end{aligned}$$

where Re_i is the i th entry in (17).

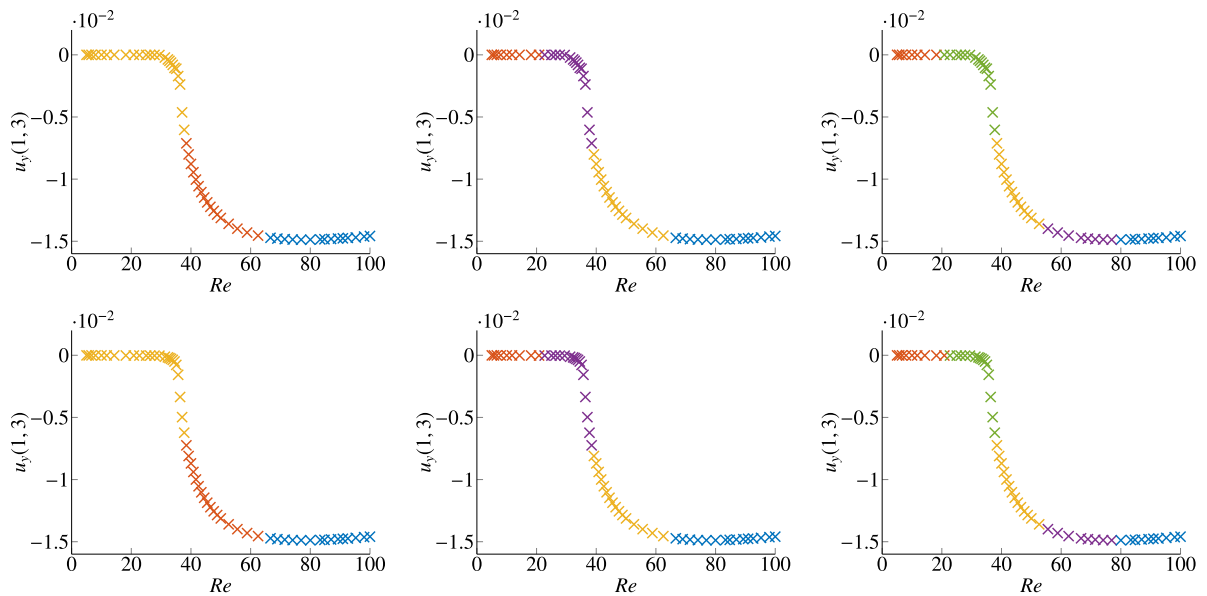


Fig. 6. Left to right: k-means clusterings of snapshots for 3, 4, and 5 clusters correspond to the use of the SEM (top row) and FEM (bottom row) methods.

3.2.3. Results for the online stage: the use of the ROM bases with local bases

Not surprisingly given the close agreement between the snapshot results obtained using the two spatial discretization methods, their application to the use of ROMs also result in very close agreements. As described in Section 2.4, we use the clustering of the snapshots to produce a POD basis for each cluster. Having those bases in hand, we then use them to determine ROM solutions, i.e., solutions of (7) using local POD bases. Of course, we do so for parameter values, i.e., Reynolds numbers in the current setting, not used in the determination of snapshots. Both SEM and FEM ROM solutions were obtained using an Oseen fixed-point iteration.

Given a new parameter value, one must decide which local basis to use. In Section 2.5, two criteria were discussed for this purpose: using the mean of the clusters or using a midrange/radius approach. The difference in the approaches is that the parameter centroid takes the distribution of snapshots in parameter space into account whereas the mid-range uses only information about the extent of the cluster in parameter space. Local ROMs are computed by POD using the snapshots belonging to their respective clusters. To recover the bifurcation diagram, the parameter mean and the midrange/radius criteria are used to decide which local basis to use for a given parameter value of interest.

In Fig. 7, bifurcation diagrams and assignments of parameters to local ROM bases are given for 3, 4, and 5 clusters with respect to SEM snapshots (left to right). The rows correspond to the use of the distance to parameter mean criteria (top row) and using the distance to parameter mid-range and cluster radius criteria (bottom row). The crosses correspond to full-order solutions for parameter values not used to generate the snapshots; they are color coded to indicate which cluster they belong to. The solid curves correspond to ROM solutions using many more new parameter values with the color coding corresponding to which cluster, and therefore which local ROM basis, a new parameter is assigned to. We observe that the midrange/radius criterion assigns new parameters to the correct local basis. On the other hand, mis-assignments of local bases occur when the parameter mean criteria is used. The same information for local ROMs relative to FEM discretization are given in Fig. 8. We note that the results are in an excellent agreement with those in Fig. 7.

Errors with respect to full-order models and comparisons to global ROM models. The goal of reduced order modeling is to obtain accurate approximations of full-order model solutions while incurring significantly reduced compute costs. In addition, because we propose a localized ROM basis approach based on clustering, comparison of the accuracy and compute times for using local ROMs should also be made with respect to a global ROM.

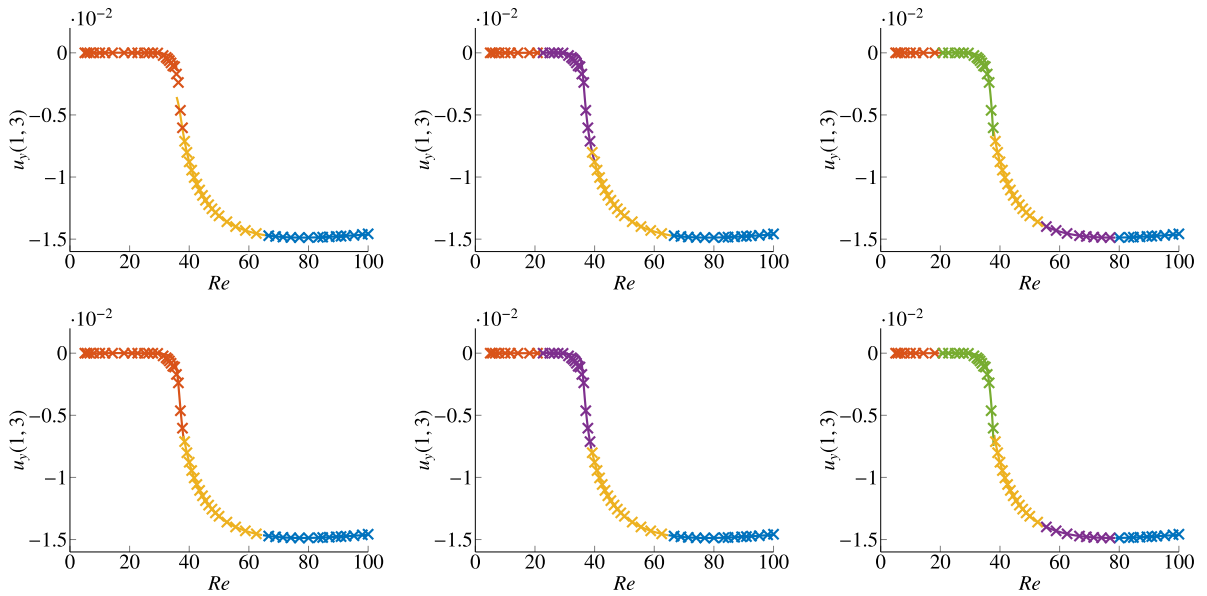


Fig. 7. For the SEM spatial discretization and for 3, 4, and 5 clusters (left to right), bifurcation diagrams and assignments of parameters to local ROMs using the distance to parameter mean criterion (top row) and using the distance to parameter mid-range and cluster radius criterion (bottom row). The crosses correspond to full-order solutions for parameter values not used to generate the snapshots whereas the solid curves correspond to ROM solutions.




Fig. 8. The same information as in Fig. 7 for the FEM spatial discretization.

Quantitative comparisons are made with respect to the relative error

$$error = \frac{\|u_{full-order} - u_{rom}\|_{L^2(\Omega)}}{\|u_{full-order}\|_{L^2(\Omega)}}, \tag{18}$$

where u_{rom} could correspond to local or global ROMs. An investigation of ROM stability and error certification with respect to the Navier–Stokes equations can be found in [49].

Table 1

For the jet problem: cluster accuracy and basis size comparisons.

3 clusters							
ROM	$Re=6$	$Re=16.333$	$Re=26.666$	$Re=37$	$Re=47.333$	$Re=56.666$	
Global-1	0.00054	0.00052	0.00052	0.0038	0.00079	0.945 (n.c.)	
Global-2	0.0027	0.0016	0.00068	0.0043	0.00091	0.00099	
Local	0.00081	0.00171	0.00121	0.011/0.021	0.00081	0.0020	
ROM	$Re=68$	$Re=78.333$	$Re=88.666$	$Re=99$	mean	max	basis size
Global-1	1.02 (n.c.)	1.00 (n.c.)	0.996 (n.c.)	0.995 (n.c.)	0.50	1.02	39
Global-2	0.0013	0.14 (n.c.)	0.26 (n.c.)	0.67 (n.c.)	0.11	0.67	21
Local ROM	0.0013	0.0084	0.0050	0.0167	0.0064	0.0167	11/7/21
5 clusters; *=10/8/5/14/17							
ROM	$Re=6$	$Re=16.333$	$Re=26.666$	$Re=37$	$Re=47.333$	$Re=56.666$	
Global-1	0.00054	0.00052	0.00052	0.0037	0.00079	0.00078	
Global-2	0.0043	0.02	0.0082	0.0064	0.0017	0.0019	
Local	0.00054	0.00052	0.00060	0.0285	0.015	0.00078	
ROM	$Re=68$	$Re=78.333$	$Re=88.666$	$Re=99$	mean	max	basis size
Global-1	0.1415 (n.c)	0.240 (n.c)	0.611 (n.c)	0.829 (n.c)	0.18	0.829	54
Global-2	0.0022	0.178 (n.c)	1.64 (n.c)	1.08 (n.c)	0.29	1.64	17
Local	0.0013	0.0084	0.005	0.045	0.011	0.045	*
7 clusters; **=8/6/4/4/8/6/16							
ROM	$Re=6$	$Re=16.333$	$Re=26.666$	$Re=37$	$Re=47.333$	$Re=56.666$	
Global-1	0.00054	0.00052	0.00052	0.0037	0.00079	0.00079	
Global-2	0.016	0.0077	0.0052	0.018	0.0039	0.0068	
Local	0.00054	0.00052	0.00057	0.0036	0.0017	0.0014	
ROM	$Re=68$	$Re=78.333$	$Re=88.666$	$Re=99$	mean	max	basis size
Global-1	0.152 (n.c)	0.23 (n.c)	0.619 (n.c)	0.712 (n.c)	0.17	0.712	52
Global-2	0.41 (n.c.)	0.48 (n.c)	7.8 (n.c)	1.7 (n.c)	1.04	7.8	16
Local	0.0013	0.0083	0.0050	0.017	0.004	0.017	**

We note that the offline computational cost is dominated by constructing the discrete trilinear form appearing in the weak form of the Navier–Stokes equations and used to project the nonlinearity; that cost grows cubically in the reduced basis size L . The online computation costs are negligible by comparison, because no operations involving the full-order model dimension are necessary.

Table 1 provides, for 3, 5, and 7 clusters, the relative error between the full-order solutions and the local and global ROM solutions as well as the size of the bases. Ten values of the Reynolds number are used to determine solutions of the full-order model, the local ROM, and the two global ROMs. Although lying within the Reynolds number range of interest, the Reynolds number values chosen are all different from those used to generate the snapshots used to construct the ROMs. One global ROM, referred to as Global-1 in the table, uses a basis of dimension equal to the sum of the dimensions of all the local ROM sizes. Whereas this global ROM has a larger compute time than the local ROM approach, it can serve as a baseline reference for assessing the approximation quality that local ROMs provide. The second global ROM, referred to as Global-2 in the table, uses a basis of dimension equal to the maximum dimension of the local bases used so that it has a similar offline computational time as that of the local ROM approach. For the local bases, results are given for both the use of the cluster mean and midrange/radius criteria when selecting what local basis to use. Values marked with (n.c.) correspond to non-converged cases; the iteration is terminated after 3000 iterates, whereas convergence typically happens after about 100 iterations. In this case, we report the relative error value for last iterate computed. Also, the arithmetic mean and maximum of relative errors (non-converged values are included) over the 10 realizations are also given in the tables.

For the 3-cluster case, the local POD basis sizes are 11, 7, and 21, respectively, with 11 corresponding the smaller values of Re and 21 to the larger values. The two global basis dimensions are thus 39 and 21. The switch between local ROMs occurs at $Re = 35$ and $Re = 65$ for the cluster mean criterion and at $Re = 38$ and $Re = 64$ for the midrange/radius criterion. Thus, two error values are shown for $Re = 37$. For the 5-cluster case, the switches

between local ROMs occur at $Re = 21, 38, 56,$ and 78 for the parameter cluster mean criterion and $Re = 19.5, 38.1, 54.1,$ and $Re = 78.5$ for the cluster midrange/radius criterion. For the 7-cluster case, the switches between local ROMs occurs at $Re = 20.2, 34.9, 42.7, 51.7, 65.2,$ and 82.1 for the parameter cluster mean criterion and at $Re = 19.5, 36.7, 42.1, 51.3, 64.6,$ and 81.7 for the cluster midrange/radius criterion.

The tabulated results suggest that the use of local reduced bases provide a more efficient and usually more accurate ROM approach when compared to the use of global reduced bases. For the one case in which the two different local basis selection criteria differed (see Table 1), the parameter cluster midrange/radius criterion resulted in a smaller error.

3.3. Example 2: Rayleigh–Bénard cavity flow

As a second example, we consider Rayleigh–Bénard (RB) cavity flow, introduced in [50] and widely studied since then. This test case features a rich bifurcating behavior despite the very simple geometrical setting. The domain Ω is a rectangular cavity with height 1 and length A . We choose the bottom left corner of the cavity as the origin of the coordinate system. The vertical cavity walls are maintained at constant temperatures T_0 (left wall) and $T_0 + \Delta T$ (right wall) with $\Delta T > 0$, whereas the horizontal walls are thermally insulated. A convective flow is induced by the horizontal component of the temperature gradient. The fluid is assumed to have density ρ and viscosity ν at the reference temperature T_0 and to satisfy homogeneous boundary condition on the cavity walls.

The strong formulation of the RB cavity problem, as stated in [50], is given by (12) with $\mathbf{f} = (0, g\beta(T - T_0)(x/A)\mathbf{e}_y)^T = (0, g\beta\Delta T(x/A)\mathbf{e}_y)^T$. Here, g denotes the magnitude of the gravitational acceleration, β is the coefficient of thermal expansion, x is the horizontal coordinate, and \mathbf{e}_y is the unit vector directed along the vertical axis. The Grashof number

$$Gr = \frac{g\beta\Delta T}{Av^2} \quad (19)$$

can be used to characterize the flow regime. This definition already takes into account the fact that the height of the cavity is 1. The Grashof number can be thought of as the ratio of buoyancy forces to viscous forces. For large Grashof numbers buoyancy forces are dominant over viscous forces and vice versa. Note that with (19) we can write $\mathbf{f} = (0, Gr\nu^2x)^T$. The Prandtl number for this problem is zero. Most of the numerous numerical studies of this problem consider $A = 4$ as was done in [50].

Depending on the governing parameters and the initial state, three branches of steady state solutions exist with one, two, or three primary circulations, also referred to as rolls. If the initial state is fluid at rest, for sufficiently small value of Gr , a steady single roll flow is observed. For a value of Gr greater than a first critical value, the steady-state velocity field features two rolls. The two-roll flow is a stable configuration until Gr exceeds a second critical value, at which the steady-state flow has three rolls. Snapshots of the flow fields for three stable scenarios are illustrated in Section 3.3.1. The passage of Gr from one subregion of the bifurcation diagram to another causes a *discontinuous change* in that the number of rolls changes. Note that all steady-state flows are centrally symmetric, i.e., with respect to rotation through 180° about the vertical centerline of the cavity, and also note that all rolls rotate in a counter-clockwise direction.

The model is defined on a rectangular domain $\Omega = [0, 4] \times [0, 1]$, i.e., the aspect ratio is set to $A = 4$. The Grashof number Gr is considered as a parameter ranging over the interval $[10^3, 10^5]$; the viscosity ν is set to one.

Because of the close agreement between the results obtained using the SEM and FEM spatial discretizations for the jet problem, for the sake of brevity, for the cavity problem we only report on numerical results obtained for the SEM discretization. The cavity domain is subdivided into 24 quadrilateral elements as shown in Fig. 9 and modal Legendre ansatz functions of order 16 are used in every element and solution component. This leads to 6321 global degrees of freedom in each component, i.e., horizontal velocity, vertical velocity, and pressure for the time-dependent simulation.

Timings. The full-order solutions are obtained using time-marching scheme with the C++ software Nektar++. On the other hand, the reduced-order solutions are obtained using a fixed point iteration scheme implemented in python. As a result, a direct comparison of timings is not possible. However, comparing a single step of the fixed-point scheme using for the full-order problem with the online reduced-order problem, compute times reduce from 100 s to 0.1 s.




Fig. 9. The 24 spectral elements used in the simulations, resulting in 18,963 degrees of freedom.




Fig. 10. Steady state solutions of the cavity flow problem for Grashof numbers $20 \cdot 10^3$, $40 \cdot 10^3$, and $98 \cdot 10^3$ (top to bottom).

In general, the online runtime is determined by the reduced basis size, since dense linear systems of this size need to be solved and the parameter-dependent system matrix is formed by evaluating the affine parameter dependence. Eliminating for a single iteration step all calls to the high-order discretization size, the compute times are 0.018s for basis dimension 10 and 0.075s for basis dimension 40. However, when computing for more than 1000 parameter locations with more than 1000 fixed-point iterations each, the ROM compute time becomes important and the local ROMs outperform the global ROMs by a factor of 1.5 to 3.

The overhead generated by the cluster selection strategy is negligible since the number of clusters K is small. For a new parameter, K euclidean distances need to be computed to find the closest local cluster, but this is done once per parameter since the same cluster is used for all fixed-point iterations.

3.3.1. Results for the offline stage: the construction of local ROM bases

In this section, we provide numerical results illustrating the different aspects of local ROM construction process for the cavity problem.

Snapshots. Fig. 10 shows representative steady state solutions of this model for Grashof number $Gr = 20 \cdot 10^3$, $40 \cdot 10^3$, and $98 \cdot 10^3$, respectively. For Gr up to $30 \cdot 10^3$, steady state solutions feature a single roll, whereas two rolls develop after that. For Gr in the interval between $92 \cdot 10^3$ and $98 \cdot 10^3$, three rolls are observed.

There also exist time-periodic solutions. For example, for $Gr = 40 \cdot 10^3$, in addition to the steady state solution with two rolls, a time-periodic solution with one roll exists. Depending on the initial condition, either solution can be obtained. For $Gr = 100 \cdot 10^3$ a time-periodic solution with three rolls also exists.

To compute a steady state, the initial condition is either given as zero or a steady state at a close parameter location. The time-step size is set to 10^{-7} and about 10^8 time steps are required to reach a steady state. Steady state




Fig. 11. Bifurcation diagram for the cavity problem corresponding to the horizontal velocity component at (0.7, 0.7).

solutions with one roll were found for $Gr \in [10 \cdot 10^3, 22 \cdot 10^3]$, two rolls for $Gr \in [30 \cdot 10^3, 90 \cdot 10^3]$, and three rolls for $Gr \in [92 \cdot 10^3, 98 \cdot 10^3]$.

Bifurcation diagram. Using full-order solutions, e.g., snapshots, we have the bifurcation diagram corresponding to the point value of the horizontal component of the velocity at (0.7, 0.7) shown in Fig. 11. The two bifurcations are clearly evident, although the one at lower Gr number is “less severe” than the one at higher Gr number. The snapshot locations are at the following Grashof numbers:

$$Gr = \{10 \cdot 10^3, 15 \cdot 10^3, 17 \cdot 10^3, 20 \cdot 10^3, 22 \cdot 10^3, 30 \cdot 10^3, 35 \cdot 10^3, \\ 40 \cdot 10^3, 45 \cdot 10^3, 50 \cdot 10^3, 55 \cdot 10^3, 60 \cdot 10^3, 65 \cdot 10^3, 70 \cdot 10^3, \\ 80 \cdot 10^3, 90 \cdot 10^3, 92 \cdot 10^3, 93 \cdot 10^3, 94 \cdot 10^3, 95 \cdot 10^3, 96 \cdot 10^3, 98 \cdot 10^3\}.$$

The sampling is finer between 92 and 98 to have a few solutions with three rolls in the snapshot space.

Clustering. In Fig. 12, the k-means variance is plotted vs. the number of clusters in the range 2 to 10. We observe a faster reduction in the variance compared to what we saw for the jet problem; compare with Fig. 5. The clustering algorithm separates the one, two, and three roll snapshots from each other. Using four clusters, the two-roll snapshots are separated into two clusters; for five clusters, the one-roll snapshots are additionally separated into two clusters. This information can be gleaned from Fig. 13. There, the crosses correspond to full-order solutions color coded to show which k-means snapshot cluster mean it is closest to, i.e., the colors correspond the different k-means clusters and therefore to different local ROM bases.

The clustering with 3 clusters is separated as follows:

cluster 1 of 3: $Gr = \{30 \cdot 10^3, 35 \cdot 10^3, 40 \cdot 10^3, 45 \cdot 10^3, 50 \cdot 10^3, 55 \cdot 10^3, 60 \cdot 10^3, 65 \cdot 10^3, 70 \cdot 10^3, 80 \cdot 10^3, 90 \cdot 10^3\}$
 cluster 2 of 3: $Gr = \{10 \cdot 10^3, 15 \cdot 10^3, 17 \cdot 10^3, 20 \cdot 10^3, 22 \cdot 10^3\}$ cluster 3 of 3: $Gr = \{92 \cdot 10^3, 93 \cdot 10^3, 94 \cdot 10^3, 95 \cdot 10^3, 96 \cdot 10^3, 98 \cdot 10^3\}$.

The clustering with 4 clusters is separated as follows:

cluster 1 of 4: $Gr = \{10 \cdot 10^3, 15 \cdot 10^3, 17 \cdot 10^3, 20 \cdot 10^3, 22 \cdot 10^3\}$ cluster 2 of 4: $Gr = \{60 \cdot 10^3, 65 \cdot 10^3, 70 \cdot 10^3, 80 \cdot 10^3, 90 \cdot 10^3\}$ cluster 3 of 4: $Gr = \{30 \cdot 10^3, 35 \cdot 10^3, 40 \cdot 10^3, 45 \cdot 10^3, 50 \cdot 10^3, 55 \cdot 10^3\}$ cluster 4 of 4: $Gr = \{92 \cdot 10^3, 93 \cdot 10^3, 94 \cdot 10^3, 95 \cdot 10^3, 96 \cdot 10^3, 98 \cdot 10^3\}$.

The clustering with 5 clusters is separated as follows :

cluster 1 of 5: $Gr = \{30 \cdot 10^3, 35 \cdot 10^3, 40 \cdot 10^3, 45 \cdot 10^3, 50 \cdot 10^3\}$ cluster 2 of 5: $Gr = \{10 \cdot 10^3, 15 \cdot 10^3, 17 \cdot 10^3, 20 \cdot 10^3, 22 \cdot 10^3\}$ cluster 3 of 5: $Gr = \{92 \cdot 10^3, 93 \cdot 10^3, 94 \cdot 10^3, 95 \cdot 10^3, 96 \cdot 10^3, 98 \cdot 10^3\}$ cluster 4 of 5: $Gr = \{80 \cdot 10^3, 90 \cdot 10^3\}$ cluster 5 of 5: $Gr = \{55 \cdot 10^3, 60 \cdot 10^3, 65 \cdot 10^3, 70 \cdot 10^3\}$.

Assigning new parameters to a local basis. As before, we compare two approaches to the assignment problem: the distance to the parameter cluster mean and the distance to mid-range/cluster radius criteria. For both criteria, we then determine the bifurcation diagram using the local ROMs; see Fig. 13. As mentioned above, the crosses in that figure correspond to a few full-order solutions; the solid curves are determined from 100 local ROM solutions.

From the left-column plots in Fig. 13, we observe that the parameter cluster mean criterium mis-assigns new parameters in the range $Gr \in [80 \cdot 10^3, 90 \cdot 10^3]$ to a three-roll basis instead of the appropriate two-roll basis, resulting in a large approximation error. In the three cluster case, the local ROM of one roll solutions is also mistakenly used




Fig. 12. k-means variance versus number of clusters for the cavity problem.

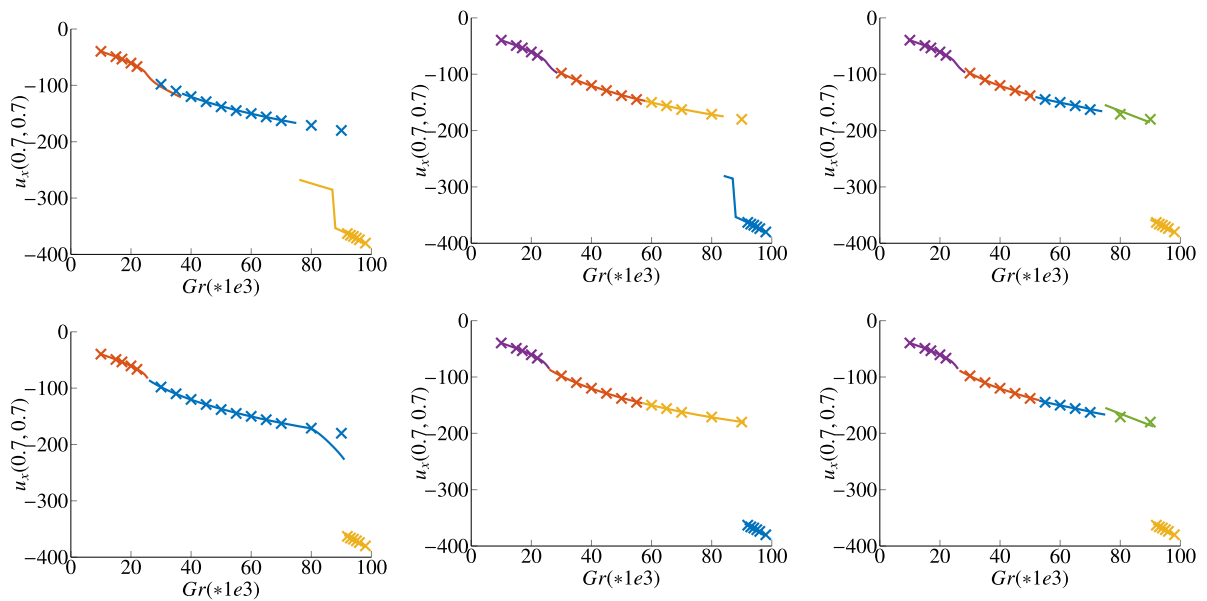


Fig. 13. Recovered bifurcation diagram with 3 (left), 4 (middle), and 5 (right) clusters using the distance to parameter mean criterion (first row) and using the distance to parameter mid-range and cluster radius criterion (second row).

for two roll solutions corresponding to $Gr \in [25 \cdot 10^3, 35 \cdot 10^3]$. Other than these mis-assignments the bifurcation diagram is recovered accurately.

From the right-column plots in Fig. 13, we observe that the midrange/radius criterion does not make such mis-assignments and only visibly deviates from the correct bifurcation diagram for the three cluster case in the range $Gr \in [80 \cdot 10^3, 90 \cdot 10^3]$. The likely reason is that the problem has not been sampled densely enough in that range.

Errors and local basis sizes. For the cavity, 5 reference solutions were computed. Table 2 shows the relative errors and basis sizes for 2, 3, 5, and 7 clusters, giving the same information as the corresponding tables for the jet problem. The local basis sizes are ordered as they are applied with increasing Grashof number.

For two clusters, the switching between local ROMs occurs at $Gr = 91.0 \cdot 10^3$ according to the midrange criterion and $Gr = 69.3 \cdot 10^3$ according to the parameter mean criterion. For three clusters, it occurs at $Gr = 36.6 \cdot 10^3$ and $Gr = 75.5 \cdot 10^3$ according to the parameter mean criterion. According to the midrange criterion, switching occurs at $Gr = 26.0 \cdot 10^3$ and $Gr = 91.0 \cdot 10^3$. For five clusters, it occurs at $Gr = 16.1 \cdot 10^3, 29.8 \cdot 10^3, 55.0 \cdot 10^3, 82.3 \cdot 10^3$ according to the parameter mean criterion. According to the midrange criterion, switching occurs at $Gr = 16.0 \cdot 10^3, 26.0 \cdot 10^3, 52.5 \cdot 10^3, 91.0 \cdot 10^3$. For seven clusters, the switching between local ROMs occurs at $Gr = 14.3 \cdot 10^3, 26.8 \cdot 10^3, 42.5 \cdot 10^3, 57.5 \cdot 10^3, 75.0 \cdot 10^3, 89.8 \cdot 10^3$ according to the parameter mean criterion. According to the midrange criterion, switching occurs at $Gr = 12.5 \cdot 10^3, 26.0 \cdot 10^3, 42.5 \cdot 10^3, 57.5 \cdot 10^3, 75.0 \cdot 10^3, 91.0 \cdot 10^3$.

Table 2

For the cavity problem: cluster accuracy and basis size comparisons.

2 clusters								
ROM	$Gr=12e3$	$Gr=32e3$	$Gr=52e3$	$Gr=75e3$	$Gr=97e3$	mean	max	basis size
Global-1	0.0027	0.0036	0.547	0.667	0.828	0.41	0.828	18
Global-2	0.0029	0.0024	0.114	1.34	0.80	0.45	1.34	15
Local	0.0028	0.0017	0.0017	0.0015/2.15	0.00018	0.36	2.15	15/3
3 clusters								
ROM	$Gr=12e3$	$Gr=32e3$	$Gr=52e3$	$Gr=75e3$	$Gr=97e3$	mean	max	basis size
Global-1	0.0027	0.0036	0.547	0.667	0.828	0.41	0.828	18
Global-2	0.0021	0.0035	0.0036	0.49	1.03	0.31	1.03	11
Local	0.0028	2.177 / 0.0017	0.0017	0.0015	0.00018	0.36	2.177	4/11/3
5 clusters: * = 2/2/3/6/3								
ROM	$Gr=12e3$	$Gr=32e3$	$Gr=52e3$	$Gr=75e3$	$Gr=97e3$	mean	max	basis size
Global-1	0.0024	0.0027	0.722	0.983	1.086	0.56	1.086	16
Global-2	0.131	0.0385	0.0114	0.007	0.0004	0.038	0.131	6
Local	0.0056	0.0055	0.5365	0.0031	0.00018	0.11	0.5365	*
7 clusters: ** = 1/3/2/2/2/2/3								
ROM	$Gr=12e3$	$Gr=32e3$	$Gr=52e3$	$Gr=75e3$	$Gr=97e3$	mean	max	basis size
Global-1	0.0029	0.0024	0.114	1.34	0.80	0.45	1.34	15
Global-2	0.255	2.83	0.967	2.41	0.238	1.34	2.83	3
Local	0.0587	0.0050	0.0090	0.013/0.0049	0.00018	0.015	0.0587	**

Because reference value $Gr = 75.0 \cdot 10^3$ corresponds to a switching point, results for both local ROM approximations are shown.

3.4. Further discussion of SEM errors for both examples

A mean accuracy of 2% was achieved by local ROMs in 5 out of 8 test cases when using the parameter mean criterion and in 7 out of 8 test cases when using the midrange rule, including the 2 and 3 cluster cases for the second example. A mean accuracy of 1% was achieved by local ROMs in 3 out of 8 test cases, when using the parameter mean criterion and in 5 out of 8 test cases when using the midrange rule, again including the 2 and 3 cluster cases for the second example.

What goes wrong in the failing cases?

First, for the 3 cluster jet case, there is a large error of 2.8% at $Re = 37$. This value is close to a switching point (at $Re = 38$) and the approximation quality tends to degrade at such points as there is a transition period in which one local ROM leaves its parameter range of validity, whereas the parameter range of validity of the next local ROM has not yet started. Also a large error of 4.5% occurs at $Re = 99$. This is the highest Re value considered and all ROMs have difficulties recovering solutions accurately as the Reynolds number approaches 100. Remedies for this can be to further increase the local ROM size or add more high Re snapshots to the sampling data.

Second, for the 5 cluster cavity case, there is a large error of 53% at $Gr = 52.0 \cdot 10^3$. This value is also close to a switching point which is at $Gr = 52.5 \cdot 10^3$ and $Gr = 55 \cdot 10^3$, respectively, and the same problem of being “in between” two clusters holds.

Third, for the 7 cluster cavity case, there is a large error of 5.9% at $Gr = 12.0 \cdot 10^3$. In this case, the local ROM has only one basis function, which does not provide good approximations. A remedy for this would be a higher sampling density, such that ROMs having only one basis function are avoided.

By contrast, the best mean accuracy a global ROM achieved is 3.8% for the 5 cluster cavity case. All other global ROMs are above 10% mean accuracy. This means that no global ROM could accurately recover all solutions in the parameter domain of interest.

For the jet flow, the sample density is higher around the bifurcation point at approximately $Re = 33$. An increasing global basis size improves approximation accuracy around the bifurcation point at the cost of accuracy at higher values of Re . Each global ROM for the jet flow failed to recover solutions for Re higher than 88. Because




Fig. 14. Steady state solution for kinematic viscosity $\nu = 0.1$ and orifice width $w = 0.5$.

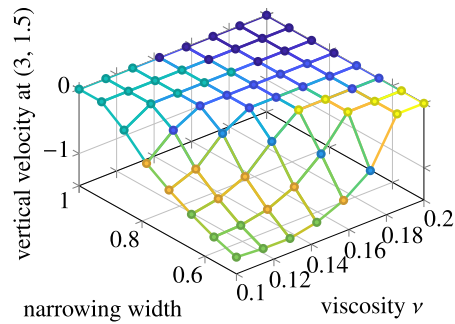


Fig. 15. Steady state vertical component of the velocity computed by SEM and evaluated at the point $(3.0, 1.5)$ for the 2-parameter study. A different marker color is used for each of the 7 clusters.

flow behavior changes more significantly at higher Re , a global ROM would need to incorporate that into its projection space. Local ROMs, on the other hand, do recover the solutions for Re higher than 88.

In the cavity example, the fixed point algorithm does converge for the global ROMs, but solutions for Grashof numbers higher than $Gr = 52.0 \cdot 10^3$ are completely inaccurate. The clustering always separates solutions with 3 rolls from solutions with 1 and 2 rolls, which makes the local ROM approach work. In the global ROMs, solutions with 1, 2, and 3 rolls are all represented in the same projection space, which could explain the inaccurate approximations.

Since the snapshot sampling is not very dense, one can only expect an approximation accuracy of three to four digits, unless solving exactly at a snapshot location. There are a number of cases where the accuracy of both global and the local ROM is below 0.5%. In such a case, we consider the solution as well-recovered by all the different ROMs. The reason why one method shows 0.1% error and the other 0.05% may be due to the larger number of basis functions available to the global ROM.

To conclude, a local ROM approach can work where global ROMs fail. It is important to be aware that the approximation quality can degrade in transition regions between two clusters and that, in those regions, the full-order model has to be sampled fine enough.

3.5. A two parameter study

To illustrate the feasibility of our approach in the case of a problem with multiple parameters, we consider jet flow with variable kinematic viscosity ν and emanating from an orifice of variable width w , symmetrically located with respect to the channel centerline. We let ν vary in the range $[0.1, 0.2]$ and w in the range $[0.5, 1.0]$. As an example, Fig. 14 shows the steady state solution for $\nu = 0.1$ and $w = 0.5$.

While for the one parameter study in Section 3.2 we have assumed a priori knowledge of the bifurcation, here we drop this assumption. In principle, our approach works well as long as the snapshot sampling is sufficiently fine. Like in Section 3.2, the vertical component of the velocity at the point $(3.0, 1.5)$ computed with SEM is plotted in the bifurcation diagram reported in Fig. 15. Snapshots have been computed on a uniform 8×8 grid and have been divided into 7 clusters. In Fig. 15 a different marker color is used for each of the 7 clusters.

The local ROMs are evaluated over a fine uniform grid of 30×30 snapshot locations. We consider three criteria to associate a new parameter location with a local ROM basis. The first criterion is an extension of the “mean of clusters” criterion discussed in Section 2.5: it entails finding the closest parameter centroid and using

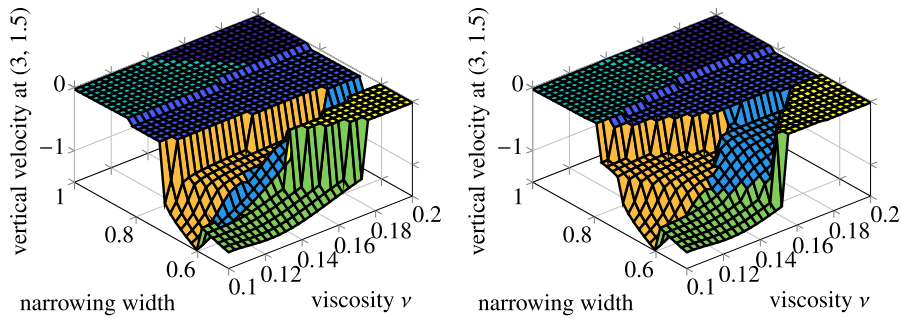


Fig. 16. Bifurcation diagram and assignment of parameters to local ROMs using the closest parameter centroid criterion (left) and the weighted closest parameter centroid (right).




Fig. 17. Bifurcation diagram and assignment of parameters to local ROMs using the closest snapshot location criterion.

the corresponding local ROM basis. The second criterion is a slight modification of the first: it uses a weight to assign the closest parameter centroid and the corresponding local ROM basis. The weighting is given by the inverse interval length for each parameter. So, the viscosity direction is weighted five times more than the geometry direction since the interval is only one fifth in length. The bifurcation diagrams obtained with these two criteria are shown in Fig. 16. The distance to parameter centroid criterion does not manage to recover the bifurcation diagram well (see Fig. 16 (left)), which seems largely due to the local ROM assignment scheme. This can be alleviated by using the weighted distance to parameter centroid as shown in the right panel of Fig. 16. However, even with this second criterion the local ROM bifurcation diagram does not match well the FOM bifurcation diagram.

In order to improve the results in Fig. 16, we consider a third criterion: we find that the closest snapshot location and consider the local ROM basis that includes this snapshot. The corresponding bifurcation diagram, reported in Fig. 17, is the most accurate reconstruction of the FOM bifurcation diagram in Fig. 15. However, this third criterion comes with an additional computational time if an unstructured grid is used for the snapshot generation.

The preliminary results reported in this section show that our local ROM approach works also for problems with more than one parameter. However, more research is needed to find an accurate and inexpensive criterion to assign the local ROM basis. This will be investigated more thoroughly in a follow-up paper.

4. Concluding remarks

In the paper, a k-means clustering of snapshots is used as the starting point for constructing a ROM that uses localized bases to treat PDE problems having bifurcating solutions. A recipe for detecting which local basis to use for any given parameter point not used to determine the snapshots is also given. Careful attention given to account for the differences between bifurcations that cause continuous or discontinuous changes in the solution. At this point, for the sake of simplicity, the new methodology is applied to problems for which the bifurcation diagram is known a priori, that is, one knows which region in the bifurcation diagram a parameter belongs to. Although restrictive, there are many problems for which such knowledge is available. The efficiency of the new approach

compared to the use of global ROM bases is demonstrated for simple, one-parameter problems for both types of bifurcations.

The results obtained so far are encouraging so that current and future work focuses on further developing our approach so that it can become a useful tool for practitioners. In particular, we will focus on:

- testing the methodology for problems with several input parameters;
- automatically detecting, without having any a priori knowledge, which local basis one should use to a given parameter point;
- testing the methodology for more “realistic” problems;
- application of the methodology to control and optimization and to uncertainty quantification problems.

References

- [1] A. Noor, On making large nonlinear problems small, *Comput. Methods Appl. Mech. Engrg.* 34 (1) (1982) 955–985.
- [2] A. Noor, Recent advances and applications of reduction methods, *ASME. Appl. Mech. Rev.* 5 (47) (1994) 125–146.
- [3] A. Noor, J. Peters, Multiple-parameter reduced basis technique for bifurcation and post-buckling analyses of composite materials, *Int. J. Numer. Methods Eng.* 19 (1983) 1783–1803.
- [4] A. Noor, J. Peters, Recent advances in reduction methods for instability analysis of structures, *Comput. Struct.* 16 (1) (1983) 67–80.
- [5] F. Terragni, J. Vega, On the use of POD-based ROMs to analyze bifurcations in some dissipative systems, *Physica D* 241 (17) (2012) 1393–1405.
- [6] H. Herrero, Y. Maday, F. Pla, RB (Reduced Basis) for RB (Rayleigh–Bénard), *Comput. Methods Appl. Mech. Engrg.* 261–262 (2013) 132–141.
- [7] F. Pla, H. Herrero, J. Vega, A flexible symmetry-preserving Galerkin/POD reduced order model applied to a convective instability problem, *Comput. & Fluids* 119 (2015) 162–175.
- [8] G. Pitton, G. Rozza, On the application of reduced basis methods to bifurcation problems in incompressible fluid dynamics, *J. Sci. Comput.* 73 (1) (2017) 157–177.
- [9] G. Pitton, A. Quaini, G. Rozza, Computational reduction strategies for the detection of steady bifurcations in incompressible fluid-dynamics: Applications to coanda effect in cardiology, *J. Comput. Phys.* 344 (2017) 534–557.
- [10] A. Quaini, R. Glowinski, S. Canic, Symmetry breaking and preliminary results about a Hopf bifurcation for incompressible viscous flow in an expansion channel, *Int. J. Comput. Fluid Dyn.* 30 (1) (2016) 7–19.
- [11] K. Cliffe, E. Hall, P. Houston, E. Phipps, A. Salinger, Adaptivity and a posteriori error control for bifurcation problems I: the Bratu problem, *Commun. Comput. Phys.* 8 (4) (2010) 845–865.
- [12] K. Cliffe, E. Hall, P. Houston, E. Phipps, A. Salinger, Adaptivity and a posteriori error control for bifurcation problems III: Incompressible fluid flow in open systems with Z(2) symmetry, *J. Sci. Comput.* 47 (3) (2011) 389–418.
- [13] K. Cliffe, E. Hall, P. Houston, E. Phipps, A. Salinger, Adaptivity and a posteriori error control for bifurcation problems III: Incompressible fluid flow in open systems with O(2) symmetry, *J. Sci. Comput.* 52 (1) (2012) 153–179.
- [14] F. Pichi, G. Rozza, Reduced basis approaches for parametrized bifurcation problems held by non-linear Von Kármán equations, 2018, arXiv:1804.02014.
- [15] S.L. Brunton, J.H. Tu, I. Bright, J.N. Kutz, Compressive sensing and low-rank libraries for classification of bifurcation regimes in nonlinear dynamical systems, *SIAM J. Appl. Dyn. Syst.* 13 (4) (2014) 1716–1732, <http://dx.doi.org/10.1137/130949282>.
- [16] B. Kramer, P. Grover, P. Boufounos, S. Nabi, M. Benosman, Sparse sensing and DMD-based identification of flow regimes and bifurcations in complex flows, *SIAM J. Appl. Dyn. Syst.* 16 (2) (2017) 1164–1196, <http://dx.doi.org/10.1137/15M104565X>.
- [17] M.L. Rapun, J. Vega, Reduced order models based on local POD plus Galerkin projection, *J. Comput. Phys.* 229 (2010) 3046–3063.
- [18] D. Amsallem, B. Haasdonk, PEBL-ROM: Projection-error based local reduced-order models, *Adv. Model. Simul. Eng. Sci.* 3:6 (2016) 1–1.
- [19] D. Amsallem, M. Zahr, C. Farhat, Nonlinear model order reduction based on local reduced-order bases, *Internat. J. Numer. Methods Engrg.* 92 (2012) 891–916.
- [20] D. Amsallem, M. Zahr, K. Washabaugh, Fast local reduced basis updates for the efficient reduction of nonlinear systems with hyper-reduction, *Adv. Comput. Math. Eng. Sci.* 41 (2015) 1187–1230.
- [21] E. Kaiser, B. Noack, L. Cordier, A. Spohn, M. Segond, M. Abel., G. Daviller, J. Östh, S. Krajnović, Cluster-based reduced-order modelling of a mixing layer, *J. Fluid Mech.* 754 (2014) 365–414.
- [22] J. Östh, E. Kaiser, S. Krajnović, B. Noack, Cluster-based reduced-order modelling of the flow in the wake of a high speed train, *J. Wind Eng. Ind. Aerodyn.* 145 (2015) 327–338.
- [23] B. Peherstorfer, D. Butnaru, K. Willcox, H.-J. Bungartz, Reduced order models based on local POD plus Galerkin projection, *SIAM J. Sci. Comput.* 36 (2014) A168–A192.
- [24] Z. Zhan, W. Habashi, M. Fossati, Local reduced-order modeling and iterative sampling for parametric analyses of aero-icing problems, *AIAA J.* 53 (2015) 2174–2185.
- [25] S. Pagani, A. Manzoni, A. Quarteroni, Numerical approximation of parametrized problems in cardiac electrophysiology by a local reduced basis method, MATHICSE Technical Report 25.2017, Lausanne, 2017.
- [26] M. Ohlberger, S. Rave, F. Schindler, True error control for the localized reduced basis method for parabolic problems, in: P. Benner, M. Ohlberger, A. Patera, G. Rozza, K. Urban (Eds.), *Model Reduction of Parametrized Systems. MS&A (Modeling, Simulation and Applications)*, vol. 17, Springer, Cham, 2017.

- [27] M. Ohlberger, B. Verfurth, Localized orthogonal decomposition for two-scale Helmholtz-type problems, *AIMS Math.* 2 (Math-02-00458) (2017) 458–478.
- [28] M. Ohlberger, F. Schindler, Non-conforming localized model reduction with online enrichment: Towards optimal complexity in PDE constrained optimization, *AIMS Math.* 2 (2017) 458–478.
- [29] J. Eftang, A. Patera, E. Ronquist, An hp certified reduced basis method for parametrized parabolic partial differential equations, in: *Mathematical and Computer Modelling of Dynamical Systems*, vol. 17, Taylor & Francis, 2011, pp. 179–187.
- [30] T. Chacón Rebollo, E. Delgado Ávila, M. Gómez Mármol, S. Rubino, Assessment of self-adapting local projection-based solvers for laminar and turbulent industrial flows, *J. Math. Ind.* 8 (1) (2018).
- [31] J.H. Tu, C.W. Rowley, D.M. Luchtenburg, S.L. Brunton, J.N. Kutz, On dynamic mode decomposition: Theory and applications, *J. Comput. Dyn.* 1 (2014) 391–421.
- [32] M. Gunzburger, Finite element method for viscous incompressible flows: a guide to theory, practice and algorithms, in: *Computer Science and Scientific Computing*, Academic Press, San Diego, 1989.
- [33] A. Quarteroni, A. Valli, *Numerical Approximation of Partial Differential Equations*, Springer-Verlag, 1994.
- [34] S. Brenner, L. Scott, *The Mathematical Theory of Finite Element Methods*, Springer-Verlag, 1994.
- [35] G. Karniadakis, S. Sherwin, *Spectral/hp Element Methods for CFD*, in: *Numerical mathematics and scientific computation*, Oxford University Press, 2005.
- [36] C. Canuto, M. Hussaini, A. Quarteroni, T. Zhang, *Spectral Methods Fundamentals in Single Domains*, in: *Scientific Computation*, Springer, 2006.
- [37] C. Canuto, M. Hussaini, A. Quarteroni, T. Zhang, *Spectral Methods Evolution to Complex Geometries and Applications to Fluid Dynamics*, in: *Scientific Computation*, Springer, 2007.
- [38] J. Hesthaven, G. Rozza, B. Stamm, *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*, Springer Briefs in Mathematics, 2015.
- [39] A. Quarteroni, A. Manzoni, F. Negri, *Reduced Basis Methods for Partial Differential Equations*, in: *UNITEXT*, 92, Springer, 2016.
- [40] S. Volkwein, Proper orthogonal decomposition: Theory and reduced-order modelling, in: *Lecture Notes*, University of Konstanz, Department of Mathematics and Statistics, 2013.
- [41] R. Fearn, T. Mullin, K. Cliffe, Nonlinear flow phenomena in a symmetric sudden expansion, *J. Fluid Mech.* 211 (1990) 595–608.
- [42] D. Drikakis, Bifurcation phenomena in incompressible sudden expansion flows, *Phys. Fluids* 9 (1) (1997) 76–87.
- [43] T. Hawa, Z. Rusak, The dynamics of a laminar flow in a symmetric channel with a sudden expansion, *J. Fluid Mech.* 436 (2001) 283–320.
- [44] S. Mishra, K. Jayaraman, Asymmetric flows in planar symmetric channels with large expansion ratios, *Internat. J. Numer. Methods Fluids* 38 (2002) 945–962.
- [45] A. Fortin, M. Jardak, J. Gervais, R. Pierre, Localization of Hopf bifurcations in fluid flow problems, *Internat. J. Numer. Methods Fluids* 24 (1997) 1185–1210.
- [46] I. Sobey, P. Drazin, Bifurcations of two-dimensional channel flows, *J. Fluid Mech.* 171 (1986) 263–287.
- [47] A. Ambrosetti, G. Prodi, *A Primer of Nonlinear Analysis*, Cambridge University Press, Cambridge, 1993.
- [48] K. Goda, A multistep technique with implicit difference schemes for calculating two- or three-dimensional cavity flows, *J. Comput. Phys.* 30 (1979) 76–95.
- [49] M. Yano, A. Patera, A space-time variational approach to hydrodynamic stability theory, *Proc. R. Soc. A* 469 (2155) (2013).
- [50] B. Roux (Ed.), *Numerical Simulation of Oscillatory Convection in Low-Pr Fluids*, in: *Notes on Numerical Fluid Mechanics and Multidisciplinary Design*, 27, Springer, 1990.