NUMERICAL SOLUTION OF THE PARAMETERIZED STEADY-STATE NAVIER-STOKES EQUATIONS USING EMPIRICAL INTERPOLATION METHODS*

HOWARD C. ELMAN[†] AND VIRGINIA FORSTALL[‡]

Abstract. Reduced-order modeling is an efficient approach for solving parameterized discrete partial differential equations when the solution is needed at many parameter values. An offline step approximates the solution space and an online step utilizes this approximation, the reduced basis, to solve a smaller reduced problem at significantly lower cost, producing an accurate estimate of the solution. For nonlinear problems, however, standard methods do not achieve the desired cost savings. Empirical interpolation methods represent a modification of this methodology used for cases of nonlinear operators or nonaffine parameter dependence. These methods identify points in the discretization necessary for representing the nonlinear component of the reduced model accurately, and they incur online computational costs that are independent of N , the number of degrees of freedom of the discrete system. We will show that empirical interpolation methods can be used to significantly reduce the costs of solving parameterized versions of the Navier-Stokes equations, and that iterative solution methods can be used in place of direct methods to further reduce the costs of solving the algebraic systems arising from reduced-order models.

Key words. reduced basis, empirical interpolation, iterative methods, preconditioning

1. Introduction. Reduced-order modeling techniques are designed to obtain the numerical solution of parameterized partial differential equations (PDEs) efficiently. In settings where solutions of parameterized PDEs are required for many parameters, such as uncertainty quantification, design optimization, and sensitivity analysis, the cost of obtaining high-fidelity solutions at each parameter may be prohibitive. In this scenario, reduced-order models can often be used to keep computation costs low by projecting the model onto a space of smaller dimension with minimal loss of accuracy.

We begin with a brief statement of the *reduced basis method* for constructing a reduced-order model. Consider an algebraic system of equations $G(u) = 0$ where $u := u(\xi)$ is an unknown vector of dimension N , and ξ is vector of m input parameters. We are interested in the case where this system arises from the discretization of a PDE and N is large, as would be the case for a high-fidelity discretization. We will refer to this system as the *full model*. We would like to compute solutions for many parameters ξ . Reduced basis methods compute a (relatively) small number of full model solutions, $u(\xi_1), \dots, u(\xi_k)$, known as snapshots, and then for other parameters, $\xi \neq \xi_j$, construct approximations of $u(\xi)$ in the space spanned by $\{u(\xi_j)\}_{j=1}^k$. In the offline-online paradigm, the offline step, which may be expensive, computes the snapshots using traditional (PDE) solvers. The offline step builds a basis of the low-dimensional vector space spanned by the snapshots. The online step, which is intended to be inexpensive (because k is small), uses a projected version of the original problem (determined, for example, by a Galerkin projection) in the k -dimensional space. The projected problem, known as the *reduced model*, has a solution $\tilde{u}(\xi)$ which is an approximation of the solution $u(\xi)$.

A straightforward implementation of the reduced-basis method is only possible for linear problems that have affine dependence on the parameters. Such problems have the form $G(u) = 0$ where

$$(1.1) \quad G(u) = A(\xi)u - b = \left(\sum_{i=1}^l \varphi_i(\xi) A_i \right) u - b$$

and $\{A_i\}_{i=1}^l$ are parameter-independent matrices. Let Q be a matrix of dimensions $N \times k$ whose columns span the space spanned by the snapshots. For example, Q can be taken to be an orthogonal matrix obtained using the Gram-Schmidt process applied to $[u(\xi_1), \dots, u(\xi_k)]$; construction of Q is part of the offline step.

*This work was supported by the U. S. Department of Energy Office of Advanced Scientific Computing Research, Applied Mathematics program under Award Number DE-SC0009301, and by the U. S. National Science Foundation under grant DMS1418754.

[†]Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742 (elman@cs.umd.edu).

[‡]Applied Mathematics and Statistics, and Scientific Computation Program, Department of Mathematics, University of Maryland, College Park, MD 20742 (vhfors@gmail.com).

With this decomposition, the reduced model obtained from a Galerkin condition is $G^r(\hat{u}) = 0$ where

$$(1.2) \quad G^r(\hat{u}) = Q^T A(\xi) Q \hat{u} - Q^T b = \left(\sum_{i=1}^l \varphi_i(\xi) (Q^T A_i Q) \right) \hat{u} - Q^T b .$$

Computation of the matrices $\{Q^T A_i Q\}$ can be included as part of the offline step. With this preliminary computation, the online step requires only the summation of the terms in equation (1.2), an $O(lk^2)$ operation, and then the solution of the system of order k . Clearly, the cost of this online computation is independent of N , the dimension of the full model.

However, when this approach is applied to a nonlinear problem, the reduced model is not independent of the dimension of the full model. Consider a problem with a nonlinear component $F(u(\xi))$, so the full model is

$$(1.3) \quad G(u(\xi)) = Au(\xi) + F(u(\xi)) - b = 0 .$$

The reduced model obtained from the Galerkin projection is

$$(1.4) \quad G^r(\hat{u}(\xi)) = Q^T A Q \hat{u}(\xi) + Q^T F(Q \hat{u}(\xi)) - Q^T b = 0$$

Although the reduced operator $Q^T F(Q \hat{u}(\xi))$ is a mapping from $\mathbb{R}^k \rightarrow \mathbb{R}^k$, any nonlinear solution algorithm (e.g. Picard iteration) requires the evaluation of the operator $F(Q \hat{u}(\xi))$ as well as the multiplication by Q^T . Both computations have costs that depend on N , the dimension of the full model.

Empirical interpolation methods [4, 9, 18, 20] use interpolation to reduce the cost of the online construction for nonlinear operators or nonaffine parameter dependence. The premise of these methods is to interpolate the operator using a subset of indices from the model. The interpolation depends on an empirically derived basis that can also be constructed as part of an offline procedure. This ensures that $F(Q \hat{u}(\xi))$ is evaluated only at a relatively small number of indices. These values are used in conjunction with a separate basis constructed to approximate the nonlinear operator. The efficiency of this approach also depends on the fact that for all i , $F_i(u(\xi))$ depends on a relatively small, $O(1)$, number of components of u .

Computing the solution of the reduced model for a nonlinear operator requires a nonlinear iteration based on a linearization strategy, which requires the solution of a reduced linear system at each step. Thus, each iteration has two primary costs, the computation of the Jacobian corresponding to $Q^T F(u(\xi))$, and the solution of the linear system at each step of the nonlinear iteration. Empirical interpolation addresses the first cost, by using an approximation of $Q^T F(u(\xi))$. To address the second cost, one option is to use direct methods to solve the reduced linear systems. In [14], however, we have seen [examples where it is more expensive to solve reduced models using direct methods than to solve the original full system, because the size of the reduced model makes the cost of direct methods large. We showed in \[14\] that this difficulty can be ameliorated using preconditioned iterative methods for the reduced problem.](#) In this paper, we extend this approach, using preconditioners that are precomputed in the offline stage, to nonlinear problems solved using empirical interpolation. We explore this approach using a Picard iteration for the linearization strategy.

We will demonstrate the efficiency of combining empirical interpolation with an iterative linear solver by computing solutions of the steady-state incompressible Navier-Stokes equations with random viscosity coefficient:

$$\begin{aligned} -\nabla \cdot \nu(\cdot, \xi) \nabla \vec{u}(\cdot, \xi) + \vec{u}(\cdot, \xi) \cdot \nabla \vec{u}(\cdot, \xi) + \nabla p(\cdot, \xi) &= f & \text{in } D \times \Gamma \\ \nabla \cdot \vec{u}(\cdot, \xi) &= 0 & \text{in } D \times \Gamma \\ \vec{u}(\cdot, \xi) &= b & \text{on } \partial D \times \Gamma, \end{aligned}$$

where $\vec{u}(\cdot, \xi)$ is the flow velocity, $p(\cdot, \xi)$ is the scalar pressure, and b determines the Dirichlet boundary conditions, and the viscosity coefficient satisfies $\nu(\cdot, \xi) > 0$. Models of this type have been used to model the viscosity in multiphase flows [19, 23, 31]. The boundary data b and forcing function f could also be parameter-dependent, although we will not consider such examples here. [We note that there is an extensive literature on the use of reduced-basis and EIM methods for the Navier-Stokes equations. For example, \[24\] develops reduced-basis methods for the stream-function / vorticity formulation of the problem, \[11, 21, 33\]](#)

concern techniques and use of a posteriori error estimation for reduced solutions, [13] demonstrates the utility of EIM for models depending on a small number of physical parameters, [3] explores heuristics for simplifying inf-sup stability of reduced models, and [35] examines variants of EIM for handling nonlinearities associated with stabilization. A comprehensive set of references can be found in [26]. Our emphasis in this study is on practical assessment of solvers for reduced basis methods for this problem, and how they fare in comparison with simply solving the full system, especially when the number of parameters increases. As we will show, the costs of reduced basis methods depend on the number of parameters, and when that number is large, solver costs increase and the advantages of the reduced model are evident only when the problem resolution is high enough. In cases where reduced bases are large, however, the cost of working with them can be reduced using iterative methods rather than direct methods to handle the algebraic systems they generate, and we introduce new methods of preconditioning for such problems and demonstrate their effectiveness.

An outline of this paper is as follows. In Section 2, we outline the details of the empirical interpolation method (EIM) that we use, the so-called discrete empirical interpolation method (DEIM) [9], which is a variant of a technique developed in [4]. (See [26, Ch. 10] for discussion and comparison of different variants of this idea.) In Section 3 we introduce the steady-state Navier-Stokes equations with an uncertain viscosity coefficient and describe the full, reduced, and DEIM models for this problem. We present numerical results in Section 4 including a comparison of snapshot selection methods for DEIM and a discussion of accuracy of the DEIM. In addition, we discuss a generalization of this approach known as a gappy-POD method [8, 16, 34]. Finally, in Section 5, we discuss the use of iterative methods for solving the reduced linear systems that arise from DEIM. This includes a presentation of new preconditioning techniques for use in this setting and discussion of their effectiveness.

2. The discrete empirical interpolation method. The discrete empirical interpolation method utilizes an approximation $\tilde{F}(u)$ of a nonlinear function $F(u)$ [9]. The keys to efficiency in this algorithm are that

1. only a small number of indices of F are used to define $\tilde{F}$,
2. each component of the nonlinear function depends only on a few indices of the input variable.

The first condition is specifically addressed by empirical interpolation as given in Algorithm 1 below. For the second condition, for finite element methods, the i th component $F_i(u)$ is associated with the i th finite element basis function (say, ϕ_i); $F_i(u)$ depends only on entries u_j for which the support of ϕ_j has nontrivial intersection with the support of ϕ_i . We elaborate further on this point later in this section. The key to the accuracy of this method is to select the indices of the discrete PDE that are most important to produce an accurate representation of the nonlinear component of the solution projected on the reduced space.

Given a PDE that depends on a set of parameters $\xi = [\xi_1, \dots, \xi_m]^T$ with solution $u(\cdot, \xi)$, the *full model* is the discretized equation such that $G(u(\cdot, \xi)) = 0$. The offline step for the traditional approach to the reduced basis method takes full solutions at several parameters and constructs a matrix Q of rank k such that $\text{range}(Q) = \text{span}\{u(\cdot, \xi^{(1)}), \dots, u(\cdot, \xi^{(k)})\}$, where the solutions $u(\cdot, \xi^{(1)}), \dots, u(\cdot, \xi^{(k)})$ are known as *snapshots*. The online step approximates the true solution u with an approximation $\tilde{u} = Q\hat{u}$. Using the Galerkin projection, the reduced model, G^r , is

$$G^r(\hat{u}) = Q^T G(Q\hat{u}) .$$

and the Jacobian of $G^r(\hat{u})$ is

$$J_{G^r}(\hat{u}) = \frac{\partial G^r(\hat{u})}{\partial \hat{u}} = Q^T \frac{\partial}{\partial \hat{u}} G(\tilde{u}(\hat{u})) = Q^T \frac{\partial G}{\partial \tilde{u}} \frac{\partial \tilde{u}}{\partial \hat{u}} = Q^T J_G(\tilde{u}) Q = Q^T J_G(Q\hat{u}) Q .$$

As observed above, when the operator is linear and affinely dependent on the parameters, the online costs (of forming and solving the reduced system (1.2)) are independent of N . This is not true for nonlinear or nonaffine operators. Consider Newton's method for the reduced model with a nonlinear operator:

$$(2.1) \quad \hat{u}_{n+1} = \hat{u}_n - J_{G^r}(\hat{u}_n)^{-1} G^r(\hat{u}_n)$$

Each iteration in equation (2.1) requires the construction of $J_{G^r}(\hat{u}_n) = Q^T J_G(Q\hat{u}_n) Q$. The construction of the matrix $J_G(Q\hat{u}_n)$ as well as the multiplications by Q and Q^T have costs that depend on N .

In the DEIM approach, the snapshots $u(\xi^{(1)}), \dots, u(\xi^{(k)})$ are obtained from the full model and the reduced basis is constructed to span these snapshots. In addition, DEIM requires a separate basis to represent the nonlinear component of the solution. This basis is constructed using a matrix of snapshots of function values $S = [F(u(\xi^{(1)})), F(u(\xi^{(2)})), \dots, F(u(\xi^{(s)}))]$. **Then, an orthogonal basis is chosen to approximately span the space spanned by these snapshots.** One approach for doing this is to use a proper orthogonal decomposition (POD) of the snapshot matrix S

$$(2.2) \quad S = \bar{V} \Sigma W^T$$

where the singular values in Σ are sorted in order of decreasing magnitude and $\bar{V}$ and W are orthogonal. DEIM will use columns of $\bar{V}$ to approximate $F(\tilde{u}(\xi))$. It may happen that $k_{deim} < k$ columns of $\bar{V}$ are used here, giving a submatrix V of $\bar{V}$. We will discuss the method used to select k_{deim} in Section 4.

Given the nonlinear basis from V , the DEIM selects indices of F so the interpolated nonlinear component on the range of V in some sense represents a good approximation to the complete set of values of $F(u(\xi))$. In particular, the approximation of the nonlinear operator is

$$\bar{F}(u(\xi)) = V(P^T V)^{-1} P^T F(u(\xi))$$

where P^T extracts entries of $F(u)$ corresponding to the interpolation points from the spatial grid.¹ This approximation satisfies $P^T \bar{F} = P^T F$. To construct P , a greedy procedure is used to minimize the error compared with the full representation of $F(u)$ [9, Algorithm 1]. For each column of V , v_i , the DEIM algorithm selects the row index for which the difference between the column v_i and the approximation of v_i obtained using the DEIM model with nonlinear basis and the first $i - 1$ columns of V is maximal, that is, the index of the maximal entry of $r = v_i - \hat{V}(P^T \hat{V})^{-1} P^T v_i$, where $\hat{V}$ denotes the first $i - 1$ columns of V . We present this in Algorithm 1.

Algorithm 1 DEIM [9]

Input: $V = [v_1, \dots, v_{k_{deim}}]$, an $N \times k_{deim}$ matrix with columns made up of the left singular vectors from the POD of the nonlinear snapshot matrix S .

Output: P , extracts the indices used for the interpolation.

- 1: $\rho = \text{argmax}(|v_1|)$, the index of the maximal entry of $|v_1|$
 - 2: $\hat{V} = [v_1]$, $P = [e_\rho]$
 - 3: **for** $i = 2 : k_{deim}$ **do**
 - 4: Solve $(P^T \hat{V})c = P^T v_i$ for c
 - 5: $r = v_i - \hat{V}c$
 - 6: $\rho = \text{argmax}(|r|)$
 - 7: $\hat{V} = [\hat{V}, v_i]$, $P = [P, e_\rho]$
 - 8: **end for**
-

Incorporating this approximation into the reduced model, equation (1.4), yields

$$(2.3) \quad \bar{F}^r(\hat{u}) = Q^T \bar{F}(\tilde{u}) = Q^T V(P^T V)^{-1} P^T F(Q\hat{u}).$$

The construction of nonlinear basis matrix V and the interpolation points are part of the offline computation. Since $L^T = Q^T V(P^T V)^{-1}$ is parameter independent, it too can be computed offline. Therefore, the online computations required are to compute $P^T J_F(u)$ and assemble $L^T(P^T J_F(u))Q$. For $P^T J_F(u)$, we need only to compute the components of $J_F(u)$ that are nonzero at the interpolation points. This is where the assumption that each component of $F(u)$ (and thus $J_F(u)$) depends on only a few entries of u is utilized. With a finite element discretization, a component $F_i(u)$ depends on the components u_j for which the intersection of the support of the basis functions have nonzero measure. **There are only $O(1)$ such indices j . Consequently, every row of $J_F(u)$ contains only $O(1)$ nonzeros and for the interpolation point with index i , construction**

¹An implementation does not literally construct the matrix P ; instead, an index list is used to extract the required entries of V .

of the i th entry of $(P^T J_F(u))Q$ has $O(1)$ cost. See [2] for additional discussion of this point. The elements that must be tracked in the DEIM computations are referred to as the sample mesh. When the sample mesh is small, the computational cost of assembling $L^T(P^T J_F(u))Q$ scales not with N but with the number of interpolation points. Therefore, DEIM will decrease the online cost associated with assembling the nonlinear component of the solution. This type of strategy, where the cost of working with a discrete nonlinear function is reduced by evaluating it only at a reduced set of points, in this case the interpolation points, is known more generally as a *hyper-reduction* method; see [1, 8] for other examples of hyper-reduction strategies.

For the Navier-Stokes equations, the nonlinear component is a function of the velocity. We will discretize the velocity space using biquadratic (Q_2) elements. In this case, an entry in $F_i(u)$ depends on at most nine entries of u . Thus this nonlinearity is amenable to using DEIM. An existing finite element routine can be used for the assembly of the required entries of the Jacobian using the sample mesh, a subset of the original mesh, as the input.

The accuracy of this approximation is determined primarily by the quality of the nonlinear basis V . This can be seen by considering the error bound

$$(2.4) \quad \|F - \bar{F}\|_2 \leq \|(P^T V)^{-1}\|_2 \|(I - VV^T)F\|_2$$

which is derived and discussed in more detail in [9, Section 3.2]. There it is shown that the greedy selection of indices in Algorithm 1 limits the growth of $\|(P^T V)^{-1}\|_2$ as the dimension of V grows. The second term $\|(I - VV^T)F\|_2$ is the quantity that is determined by the quality of V . Note that if V is taken from the truncated POD of S , the matrix of nonlinear snapshots, then $\|(I - VV^T)S\|_F^2$ is minimized [2] where $\|\cdot\|_F$ is the Frobenius norm ($\|X\|_F^2 = \sum_i \sum_j |x_{ij}|^2$). So the accuracy of the DEIM approximation depends on two factors. The first is the number, k_{deim} , of singular vectors kept in the POD. The truncated matrix $V\Sigma_{deim}W_{deim}^T$ is the optimal rank- k_{deim} approximation of S , but a higher rank approximation will improve accuracy of the DEIM model. In fact the error $\|(I - VV^T)F\|_2$ approaches 0 in the limit as k_{deim} approaches N . The second factor is the quality of the nonlinear snapshots in S . The nonlinear component should be sampled well enough to capture the variations of the nonlinear component throughout the solution space. A comparison of methods for selecting the snapshot set is given in Section 4.1.

Given the full model defined in equation (1.3), let $J_G(u) = A + J_F(u)$ denote the Jacobian matrix. The Jacobian of the reduced model equation (1.4) is then

$$J_{Gr}(\hat{u}) = Q^T J_G(Q\hat{u})Q = Q^T A Q + Q^T J_F(Q\hat{u})Q,$$

and the Jacobian $\bar{F}^r(\hat{u})$ of (2.3) is

$$J_{\bar{F}^r}(\hat{u}) = Q^T J_{\bar{F}}(Q\hat{u})Q = Q^T V(P^T V)^{-1}P^T J_F(u)Q.$$

3. Steady-state Navier-Stokes equations. A discrete formulation of the steady-state Navier-Stokes equations (1.5) is to find $\vec{u}_h \in X_E^h$ and $p_h \in M^h$ such that

$$\begin{aligned} (\nu(\cdot, \xi) \nabla \vec{u}_h, \nabla \vec{v}_h) + (\vec{u}_h \cdot \nabla \vec{u}_h, \vec{v}_h) - (p_h, \nabla \cdot \vec{v}_h) &= (f, \vec{v}_h) \quad \forall \vec{v}_h \in X_0^h \\ (\nabla \cdot \vec{u}_h, q_h) &= 0 \quad \forall q_h \in M^h \end{aligned}$$

where X_E^h and M^h are finite-dimensional subspaces of the Sobolev spaces $H_0^1 = \{\vec{v} \in H^1 | \vec{v} = 0 \text{ on } \partial D\}$ and $L_2(D)$. We will use div-stable Q_2 - P_{-1} finite element (biquadratic velocities, piecewise constant discontinuous pressure). See [15, 17] for details of such discretizations. Let $\{\phi_1, \dots, \phi_{N_u}\}$ represent a basis of Q_2 and $\{\psi_1, \dots, \psi_{N_p}\}$ represent a basis of P_{-1} .

We define the following vectors and matrices, where u and p here represent the vectors of coefficients that determine $\vec{u}_h$ and p_h , respectively:

$$\mathbf{z} = \begin{bmatrix} u \\ p \end{bmatrix},$$

$$[A(\xi)]_{ij} = \int \nu(\xi) \nabla \phi_i : \nabla \phi_j, \quad [B]_{ij} = - \int \psi_i (\nabla \cdot \phi_j), \quad [C(u)]_{ij} = \int (\vec{u}_h \cdot \nabla \phi_j) \cdot \phi_i,$$

$$[\mathbf{f}]_i = (f, \phi_i), \quad [g(\vec{u})]_i = -(\nabla \cdot \vec{u}_h, \psi_i), \quad \mathbf{b}(\xi) = \begin{bmatrix} \mathbf{f} - A(\xi)u_{bc} \\ g(u_{bc}) \end{bmatrix},$$

where u_{bc} is the vector of coefficients of the discrete velocity field $\vec{u}_{bc}$ that interpolates the Dirichlet boundary data $b(\cdot, \xi)$ and is zero everywhere on the interior of the mesh. We denote the velocity solution on the interior of the mesh, $\vec{u}_{in}$, so that $\vec{u} = \vec{u}_{bc} + \vec{u}_{in}$ and $\vec{u}_{in}$ satisfies homogenous Dirichlet boundary conditions. The reduced basis is constructed using snapshots of $\vec{u}_{in}$ so the approximation of the velocity solution generated by the reduced model is of the form $\tilde{u} = u_{bc} + Q_u \hat{u}$ where the columns of Q_u correspond to a basis spanning the space of velocity snapshots with homogeneous Dirichlet boundary conditions.

3.1. Full model. With this notation, the full discrete model for the Navier-Stokes problem with parameter ξ is to find $\mathbf{z}(\xi)$ such that $G(\mathbf{z}(\xi)) = 0$ where

$$(3.1) \quad G(\mathbf{z}(\xi)) = \begin{bmatrix} A(\xi) & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} u \\ p \end{bmatrix} + \begin{bmatrix} C(u) & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} u \\ p \end{bmatrix} - \begin{bmatrix} \mathbf{f} \\ 0 \end{bmatrix}, \quad \begin{bmatrix} u \\ p \end{bmatrix} = \begin{bmatrix} u(\xi) \\ p(\xi) \end{bmatrix}.$$

We utilize a Picard iteration to solve the full model, monitoring the norm of the nonlinear residual $G(\mathbf{z}_n(\xi))$ for convergence. The nonlinear Picard iteration to solve this model is described in Algorithm 2.

Algorithm 2 Picard iteration for solving the discrete steady-state Navier-Stokes equations

1: The nonlinear iteration is initialized with the solution to a Stokes problem

$$(3.2) \quad \begin{bmatrix} A(\xi) & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} \vec{u}_{in,0} \\ p_0 \end{bmatrix} = \mathbf{b}(\xi).$$

2: Incorporate the boundary conditions

$$\vec{u}_0 = \vec{u}_{bc} + \vec{u}_{in,0}.$$

3: Solve

$$(3.3) \quad \left(\begin{bmatrix} A(\xi) & B^T \\ B & 0 \end{bmatrix} + \begin{bmatrix} C(\vec{u}_n) & 0 \\ 0 & 0 \end{bmatrix} \right) \begin{bmatrix} \delta \vec{u} \\ \delta p \end{bmatrix} = -G(\mathbf{z}_n).$$

4: Update the solutions

$$\begin{aligned} \vec{u}_{n+1} &= \vec{u}_n + \delta \vec{u} \\ p_{n+1} &= p_n + \delta p. \end{aligned}$$

5: Exit if

$$\|G(\mathbf{z}_{n+1})\|_2 < \delta \|\mathbf{b}(\xi)\|_2,$$

otherwise return to step 2.

3.2. Reduced model. Next, we present a reduced model that does not make use of the DEIM strategy. This is not meant to be practical strategy, but it is presented as a comparison to illustrate how the additional approximation used for DEIM affects both the accuracy of solutions obtained using DEIM and the speed with which they are obtained. Offline we compute a reduced basis

$$Q = \begin{bmatrix} Q_u & 0 \\ 0 & Q_p \end{bmatrix},$$

where Q_u represents the reduced basis of the velocity space and Q_p the reduced basis for the pressure space. We defer the details of this offline construction to Section 4.1. For a given Q , the Galerkin reduced model is

$$G^r(\mathbf{z}) = Q^T G(\mathbf{z}) = \left(Q^T \begin{bmatrix} A(\xi) & B^T \\ B & 0 \end{bmatrix} Q \right) \begin{bmatrix} \hat{u} \\ \hat{p} \end{bmatrix} + \left(Q^T \begin{bmatrix} C(\tilde{u}) & 0 \\ 0 & 0 \end{bmatrix} Q \right) \begin{bmatrix} \hat{u} \\ \hat{p} \end{bmatrix} - Q^T \begin{bmatrix} \mathbf{f} \\ 0 \end{bmatrix}.$$

6

Algorithm 3 Picard iteration for solving the reduced steady-state Navier-Stokes equations

- 1: Initialize the Picard iteration by solving the reduced Stokes problem

$$Q^T \begin{bmatrix} A(\xi) & B^T \\ B & 0 \end{bmatrix} Q \begin{bmatrix} \hat{u}_0 \\ \hat{p}_0 \end{bmatrix} = Q^T \mathbf{b}(\xi) .$$

- 2: Solve the reduced problem for the Picard iteration

$$\left(Q^T \begin{bmatrix} A(\xi) & B^T \\ B & 0 \end{bmatrix} Q + Q^T \begin{bmatrix} C(\tilde{u}_n) & 0 \\ 0 & 0 \end{bmatrix} Q \right) \begin{bmatrix} \delta \hat{u} \\ \delta \hat{p} \end{bmatrix} = -G^r(\tilde{\mathbf{z}}_n) .$$

Note that when the dependence on the parameters is affine, the first term in the left hand side can be computed primarily offline as in equation (1.2).

- 3: Update the reduced solutions

$$\begin{aligned} \hat{u}_{n+1} &= \hat{u}_n + \delta \hat{u} \\ \hat{p}_{n+1} &= \hat{p}_n + \delta \hat{p} . \end{aligned}$$

- 4: Update the approximation to the full solution

$$\begin{aligned} \tilde{u}_{n+1} &= \vec{u}_{bc} + Q_u \hat{u}_{n+1} \\ \tilde{p}_{n+1} &= Q_p \hat{p}_{n+1} . \end{aligned}$$

- 5: Compute $C(\tilde{u}_{n+1})$.

- 6: Compute $G^r(\tilde{\mathbf{z}}_{n+1})$.

- 7: Exit if

$$\|G^r(\tilde{\mathbf{z}}_{n+1})\|_2 < \delta \|Q^T \mathbf{b}(\xi)\|_2 ,$$

otherwise return to step 2.

Using the nonlinear Picard iteration, the reduced model is described in Algorithm 3.

After the convergence of the Picard iteration determined by $G^r(\tilde{\mathbf{z}}_{n+1})$, we compute the “full” residual $G(\tilde{\mathbf{z}}_n)$. Note that this residual is computed only once; because the cost of computing it is $O(N)$, it is not monitored during the course of the iteration. The full residual [is an easily computable quantity that](#) indicates how well the reduced model approximates the full solution, so it is used to measure the quality of the reduced model via the error indicator

$$(3.4) \quad \eta_\xi = \|G(\tilde{\mathbf{z}}_n(\xi))\|_2 / \|\mathbf{b}(\xi)\|_2 .$$

3.3. DEIM model. The DEIM model has the structure of the reduced model but with the nonlinear component F replaced by the approximation $\bar{F}$. First in the offline step, we compute V , P , and $L^T = Q_u^T V (P^T V)^{-1}$. The DEIM model is

$$(3.5) \quad G^{deim}(\mathbf{z}) = \left(Q^T \begin{bmatrix} A(\xi) & B^T \\ B & 0 \end{bmatrix} Q \right) \begin{bmatrix} \hat{u} \\ \hat{p} \end{bmatrix} + \begin{bmatrix} L^T P^T C(\tilde{u}) Q_u & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \hat{u} \\ \hat{p} \end{bmatrix} - Q^T \begin{bmatrix} \mathbf{f} \\ 0 \end{bmatrix} .$$

The computations are shown in Algorithm 4. Recall from the earlier discussion of DEIM that $P^T C(u)$ is not computed by forming the matrix $C(u)$. Instead, only the components of $C(u)$ corresponding to the indices required by P^T are constructed, using the elements of the discretization mesh that contribute to those indices. Note that error indicator η_ξ in equation (3.4) depends on $G(\tilde{\mathbf{z}})$. This quantity contains $C(\tilde{u}_n)$

and not $P^T C(\tilde{u}_n)$, so that computing it requires assembly of $C(\tilde{u}_n)$ on the entire mesh. As for the reduced model, to avoid this expense, this computation is performed only after convergence of the nonlinear iteration.

Algorithm 4 DEIM model for the steady-state Navier-Stokes equations

- 1: Initialize the Picard iteration by solving the reduced Stokes problem

$$(3.6) \quad Q^T \begin{bmatrix} A(\xi) & B^T \\ B & 0 \end{bmatrix} Q \begin{bmatrix} \hat{u}_0 \\ \hat{p}_0 \end{bmatrix} = Q^T \mathbf{b}(\xi) .$$

- 2: Solve the reduced problem for the Picard iteration

$$(3.7) \quad \left(Q^T \begin{bmatrix} A(\xi) & B^T \\ B & 0 \end{bmatrix} Q + \begin{bmatrix} L^T (P^T C(\tilde{u}_n)) Q_u & 0 \\ 0 & 0 \end{bmatrix} \right) \begin{bmatrix} \delta \hat{u} \\ \delta \hat{p} \end{bmatrix} = -G^{deim}(\tilde{\mathbf{z}}_n) .$$

Note that the term on the left can be computed cheaply as in equation (1.2) so we only need to update the upper left corner of the matrix as the Picard iteration proceeds.

- 3: Update the reduced solutions

$$\begin{aligned} \hat{u}_{n+1} &= \hat{u}_n + \delta \hat{u} \\ \hat{p}_{n+1} &= \hat{p}_n + \delta \hat{p} . \end{aligned}$$

- 4: Compute $P^T C(\tilde{u}_{n+1}) = P^T C(Q \hat{u}_{n+1})$ at required indices.
- 5: Compute $G^{deim}(\tilde{\mathbf{z}}_{n+1})$.
- 6: If

$$\|G^{deim}(\tilde{\mathbf{z}}_{n+1})\|_2 < \delta \|Q^T \mathbf{b}(\xi)\|_2 ,$$

compute the approximation to the full solution

$$\begin{aligned} \tilde{u}_{n+1} &= u_{bc} + Q_u \hat{u}_{n+1} \\ \tilde{p}_{n+1} &= Q_p \hat{p}_{n+1} . \end{aligned}$$

and exit.

Otherwise, return to step 2.

Remark. If the reduced velocity basis is $Q_u = [q_1^{(u)}, \dots, q_{k_u}^{(u)}]$, then the velocity iterate in Algorithm 3 has the form $\tilde{u}_n = \sum_{i=1}^{k_u} q_i^{(u)} \hat{u}_i$. Since the nonlinearity in the Navier-Stokes equations is only quadratic (in the convection term $(\tilde{u} \cdot \nabla) \tilde{u}$), the matrix for the linearized convection term in Algorithm 3 (step 2) is $C(\tilde{u}_n) = C(Q_u \hat{u}) = \sum_{i=1}^{k_u} C_i \hat{u}_i$, where $C_i = C(q_i^{(u)})$. This gives the reduced convection matrix

$$(3.8) \quad Q_u^T C(\tilde{u}_n) Q_u = \sum_{i=1}^{k_u} (Q_u^T C_i Q_u) \hat{u}_i .$$

As in (1.2), this matrix can be constructed at cost independent of N , the size of the full model. However, the number of terms and dense matrices needed for (3.8) is k_u rather than ℓ ($= m$ in our setting). As will be seen in Section 4, for most of the examples in this study, k_u is significantly larger than m , and the work and storage requirements for using (3.8) would be prohibitive. In contrast, the DEIM approximation $L^T P^T C(\tilde{u}_n) Q_u$ of (3.5) depends only on a single dense matrix $L^T = Q_u^T V (P^T V)^{-1}$, so its costs are minimal.

3.4. Inf-sup condition. We turn now to the construction of the reduced basis

$$Q = \begin{bmatrix} Q_u & 0 \\ 0 & Q_p \end{bmatrix}.$$

Given k_s snapshots of the full model, a natural choice is to have the following spaces generated by these snapshots,

$$(3.9) \quad \begin{aligned} \text{span}(Q_u^{snap}) &= \text{span}\{\vec{u}_{in}(\xi^{(1)}), \dots, \vec{u}_{in}(\xi^{(k_s)})\} \\ \text{span}(Q_p) &= \text{span}\{p(\xi^{(1)}), \dots, p(\xi^{(k_s)})\}. \end{aligned}$$

However, this choice of basis does not satisfy an inf-sup condition

$$\gamma_R := \min_{0 \neq q_R \in \text{span}(Q_p)} \max_{0 \neq \vec{v}_R \in \text{span}(Q_u^{snap})} \frac{(q_R, \nabla \cdot \vec{v}_R)}{|\vec{v}_R|_1 \|q_R\|_0} \geq \gamma^* > 0$$

where γ^* is independent of Q_u^{snap} and Q_p . To address this issue, we follow the enrichment procedure of [27, 28]. For $i = 1, \dots, k_s$, let $\vec{r}_h(\cdot, \xi^{(i)})$ be the solution to the Poisson problem

$$(\nabla \vec{r}_h(\cdot, \xi^{(i)}), \nabla \vec{v}_h) = (p_h(\cdot, \xi^{(i)}), \nabla \cdot \vec{v}_h) \quad \forall \vec{v}_h \in X_0^h,$$

and let Q_u^{snap} of equation (3.9) be augmented by the corresponding discrete solutions $\{\vec{r}(\xi^{(i)})\}$, giving the enriched space

$$\text{span}(Q_u) = \text{span}\{\vec{u}_{in}(\xi^{(1)}), \dots, \vec{u}_{in}(\xi^{(k_s)}), \vec{r}(\xi^{(1)}), \dots, \vec{r}(\xi^{(k_s)})\}.$$

These enriching functions satisfy

$$\vec{r}_h(\cdot, \xi^{(i)}) = \arg \sup_{\vec{v}_h \in X_0^h} \frac{(p_h(\cdot, \xi^{(i)}), \nabla \cdot \vec{v}_h)}{|\vec{v}_h|_1},$$

and thus γ_R defined for the enriched velocity space, $\text{span}(Q_u)$, together with $\text{span}(Q_p)$, satisfies the inf-sup condition

$$\gamma_R \geq \gamma_h := \min_{0 \neq q_h \in M_h} \max_{0 \neq \vec{v}_h \in X_0^h} \frac{(q_h, \nabla \cdot \vec{v}_h)}{|\vec{v}_h|_1 \|q_h\|_0}.$$

Cf. [12] for an alternative approach, in which reduced models derived from *unstable* mixed elements are treated using stabilization methods.

4. Experiments. We consider the steady-state Navier-Stokes equations (1.5) for driven cavity flow posed on a square domain $D = (-1, 1) \times (-1, 1)$. The lid, the top boundary ($y = 1$), has velocity profile

$$u_x = 1 - x^4, \quad u_y = 0,$$

and no-slip conditions $\vec{u} = (0, 0)^T$ hold on other boundaries. The source term is $f \equiv 0$. The discretization is done on a div-stable $n_{fe} \times n_{fe}$ Q_2 - P_{-1} (biquadratic velocities, discontinuous piecewise constant pressures) element grid, giving a discrete velocity space of order $N_u = (n_{fe} + 1)^2$ and pressure space of order $N_p = 3(n_{fe}/2)^2$.

To define the uncertain viscosity, we divide the domain D into $m = n_d \times n_d$ subdomains as seen in Figure 4.1, and the viscosity is taken to be constant and random on each subdomain, $\nu(\xi) = \xi_i$. The random parameter vector, $\xi = [\xi_1, \dots, \xi_m]^T \in \Gamma$, is comprised of uniform random variables such that $\xi_i \in \Gamma_i = [0.01, 1]$ for each i . Therefore, the local subdomain-dependent Reynolds number, $\mathcal{R} = 2/\nu$, will vary between 2 and 200 for this problem. Constant Reynolds numbers in this range give rise to stable steady solutions [15].

The implementation uses IFISS [30] to generate the finite element matrices for the full model. The matrices are then imported into Python and the full, reduced, and DEIM models are constructed and solved using a Python implementation run on an Intel 2.7 GHz i5 processor and 8 GB of RAM. The full model

Fig. 4.1: Flow domain with piecewise random coefficients for viscosity.

is solved with the method described in equation (4.1) using sparse direct methods implemented in the UMFPACK suite [10] for the system solves in equations (3.2) and (3.3). For this benchmark problem (with enclosed flow), these linear systems are singular [6]. This issue is addressed by augmenting the matrix, for example that of (3.2), as

$$(4.1) \quad \begin{bmatrix} A(\xi) & B^T & 0 \\ B & 0 & 0 \\ 0 & \bar{p} & 0 \end{bmatrix} \begin{bmatrix} \vec{u}_{in} \\ p \\ \lambda \end{bmatrix} = \begin{bmatrix} \mathbf{b}(\xi) \\ 0 \end{bmatrix},$$

where $\bar{p}$ is a vector corresponding to the element areas of the pressure elements. This removes the singularity by adding a constraint via a Lagrange multiplier so that the average pressure of the solution is zero [29]. The same constraint is added to the systems in equation (3.3).

4.1. Construction of Q and V . We now describe the methodology we used to compute the reduced bases, Q_u and Q_p , and the nonlinear basis V . The description of the construction of Q_u and Q_p is presented in Algorithm 5. The reduced bases Q_u and Q_p are constructed using random sampling of k_{trial} samples of Γ , denoted Γ_{trial} . The bases are constructed so that all samples $\xi \in \Gamma_{trial}$ have a residual indicator, η_ξ , less than a tolerance, τ . The procedure begins with single snapshot $\mathbf{z}(\xi^{(0)})$ where $\xi^{(0)} = E(\xi)$, the expected value of the (joint uniformly distributed) random variable ξ . The bases are initialized using this snapshot, such that $Q_u = [\vec{u}_{in,n}(\xi^{(0)}), \vec{r}(\xi^{(0)})]$ and $Q_p = [p_n(\xi^{(0)})]$. Then for each sample of Γ_{trial} , the reduced-order model is solved with the current bases Q_u and Q_p . The quality of the reduced solution produced by this reduced-order model can be evaluated using the error indicator η_ξ of (3.4). If η_ξ is smaller than the tolerance τ , the computation proceeds to the next sample. If the error indicator exceeds the tolerance, the full model is solved, and then the new snapshots, $u_{in,n}(\xi)$ and $p_n(\xi)$, and the enriched velocity $\vec{r}(\xi)$, are used to augment Q_u and Q_p . The experiments use $\tau = 10^{-4}$ and $k_{trial} = 2000$ parameters to produce bases Q_u and Q_p .

We turn now to the methodology for determining the nonlinear basis V , which is the truncated form of $\bar{V}$ defined in equation (2.2). It was shown in [9] that the choice of the nonlinear basis, V in equation (2.4), is important for the accuracy of the DEIM model. DEIM uses a POD approach for constructing the nonlinear basis. This POD has a two inputs, S , the matrix of nonlinear snapshots, and k_{deim} , the number of vectors retained after truncation. We will compare three strategies for sampling S specified in line 17 of Algorithm 5.

1. **Full(k_{trial})**. This method is most similar to the method used to generate the nonlinear basis in [9]. The matrix of nonlinear snapshots, S , is computed from the full solution at every random sample (i.e. $\{C(\vec{u}(\xi^{(i)}))\vec{u}(\xi^{(i)})\}_{i=1}^{k_{trial}}$). This computation is part of the offline step but its cost, for solving solving k_{trial} full problems, can be quite high.
2. **Full(k_s)**. This is the sampling strategy included in Algorithm 5. It saves the nonlinear component only when the full model is solved for augmenting the reduced basis, Q . Therefore the snapshot set S contains k_s snapshots.
3. **Mixed(k_{trial})**. The final approach aims to mimic the Full(k_{trial}) method with less offline work. This method generates a nonlinear snapshot for each of the k_{trial} random samples using full solutions when they are available (from full solution used for augmenting the reduced basis) and reduced solutions when they are not. As the reduced basis is constructed, when the solution to the full problem is not needed for the reduced basis (i.e. when $\eta_{\xi^{(i)}} < \tau$), use the reduced solution $\tilde{u}(\xi^{(i)})$

Algorithm 5 Construction of reduced basis Q via random sampling, construction of nonlinear basis V

Cost: k_{trial} reduced problems and k_s full problems.

- 1: $\triangleright$ Compute the reduced basis and the nonlinear snapshots
- 2: Solve the full problem $G(\mathbf{z}(\xi^{(0)})) = 0$ to tolerance δ for $\mathbf{z}(\xi^{(0)}) = \begin{bmatrix} \tilde{\mathbf{u}}(\xi^{(0)}) \\ p(\xi^{(0)}) \end{bmatrix}$.
- 3: Compute the enriched velocity, $\tilde{\mathbf{r}}(\xi^{(0)})$.
- 4: Initialize $Q_u = [\tilde{\mathbf{u}}(\xi^{(0)}), \tilde{\mathbf{r}}(\xi^{(0)})]$ and $Q_p = [p(\xi^{(0)})]$.
- 5: Save the nonlinear component of the solution $S = [C(\tilde{\mathbf{u}}(\xi^{(0)}))\tilde{\mathbf{u}}(\xi^{(0)})]$.
- 6: $k_s = 1$
- 7: **for** $i = 1 : k_{trial}$ **do**
- 8: Randomly select $\xi^{(i)}$.
- 9: Solve reduced model $G^r(\tilde{\mathbf{z}}(\xi^{(i)})) = 0$ to tolerance δ and compute the residual indicator $\eta_{\xi^{(i)}}$.
- 10: **if** $\eta_{\xi^{(i)}} > \tau$ **then**
- 11: $k_s = k_s + 1$
- 12: Solve full model $G(\mathbf{z}(\xi^{(i)})) = 0$.
- 13: Compute the enriched velocity, $\tilde{\mathbf{r}}(\xi^{(i)})$.
- 14: Add $\tilde{\mathbf{u}}(\xi^{(i)})$ and $\tilde{\mathbf{r}}(\xi^{(i)})$ to Q_u and $p(\xi^{(i)})$ to Q_p using modified Gram-Schmidt.
- 15: Add the nonlinear component to the matrix of nonlinear snapshots, $S = [S, C(\tilde{\mathbf{u}}(\xi^{(i)}))\tilde{\mathbf{u}}(\xi^{(i)})]$.
- 16: **end if**
- 17: **end for**
- 18: Compute the POD of the nonlinear snapshot matrix:

$$S = \bar{V} \begin{bmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_{k_s} \end{bmatrix} W^T.$$

- 19: Choose $k_{deim} = \text{rank}(S)$.
 - 20: Define $V = \bar{V}[:, 1 : k_{deim}]$.
 - 21: Compute P using Algorithm 1 with input V .
-

to generate the nonlinear snapshot $C(\tilde{\mathbf{u}}(\xi^{(i)}))\tilde{\mathbf{u}}(\xi^{(i)})$, where $\tilde{\mathbf{u}}(\xi^{(i)}) = \tilde{\mathbf{u}}_{bc} + Q_u \hat{\mathbf{u}}(\xi^{(i)})$ and Q_u is the basis at this value of i . Thus S contains k_{trial} snapshots, but it is constructed using only k_s full model solves.

Figure 4.2 compares the performance of the three methods for generating S when Algorithm 5 is used to generate Q_u and Q_p . For each S , we take the SVD and truncate with a varying number of vectors, k_{deim} , and plot the average of the residuals of the DEIM solution for 100 randomly generated samples of ξ . The average residual for the reduced model without DEIM is also shown. It can be seen that as k_{deim} increases, the residuals of the DEIM models approach the residual obtained without using DEIM. It is also evident that the three methods perform similarly. Thus, the Mixed(k_{trial}) approach provides accurate nonlinear snapshots with fewer full solutions than the Full(k_{trial}) method. The Full(k_s) method is essentially as effective as the others and requires both fewer full-system solves than Full(k_{trial}) and the SVD of a smaller matrix than Mixed(k_{trial}); we used Full(k_s) for the remainder of this study. For this method, $k_{deim} \leq k_s$ (in fact, in all cases $k_{deim} = k_s$), which is why its results do not fully extend across the horizontal axis in Figure 4.2; if the (slightly higher) accuracy exhibited by the other methods is needed, it could be obtained using the Mixed(k_{trial}) method at relatively little extra cost.

4.2. Online component - DEIM model versus reduced model. In Section 2, we presented analytic bounds for how accurately the DEIM approximates the nonlinear component of the model. To examine how the approximation affects the accuracy of the reduced model, we will compare the error indicators of DEIM with those obtained using the reduced model without DEIM. We perform the following offline and online computations:

1. Offline: Use Algorithm 5 with input $\tau = 10^{-4}$ and $k_{trial} = 2000$ to generate the reduced bases Q_u ,

Fig. 4.2: Comparison of methods to generate nonlinear snapshots for DEIM. DEIM residual vs. k_{deim} averaged for $n_s = 100$ samples, 32×32 finite element grid, $m = 4$ random variables, τ (tolerance for reduced basis) $= 10^{-4}$, k (rank of Q) $= 306$, k_{deim} varies. Q_u , Q_p are generated using Algorithm 5.

Q_p , V and the indices in P .

2. Online: Solve the problem using the full model, reduced model without DEIM, and the reduced model with DEIM for $n_s = 10$ sample parameter sets.

Table 4.1 presents the results of these tests. For each benchmark problem, there are four entries. Three of the entries, the first, third and fourth, show the time required to meet the stopping criterion in each of Algorithms 2 (full system), 3 (reduced system) and 4 (reduced system with DEIM) with tolerance $\delta = 10^{-8}$, together with the final relative residual of the full nonlinear system. The reduced models use a reduced residual for the stopping test; as we have observed, this is done for DEIM to make the cost of the iteration depend on the size of the reduced model rather than N , the size of the full model. (Although the reduced model without DEIM has costs that depend on N , we used the same stopping criterion in order to assess the difference between the two reduced models.) The other entries of Table 4.1, the second for each example, are the time and residual data for the full system solve with a milder tolerance, $\delta = 10^{-4}$; with this choice, the relative residual is comparable in size to that obtained for the reduced models.

For each test problem, the smallest computational time is in boldface. The times presented for the reduced and DEIM models are the CPU time spent in the online computation for the nonlinear iteration only and do not include assembly time or the time to compute the full nonlinear residual after the iteration for the reduced model has converged.

The results demonstrate the tradeoff between accuracy and time for the three models. For example, in the case of $m = 16$ and $n = 65$, the spatial dimension and parameter dimension are large enough that the DEIM model is fastest, and this is generally true for higher-resolution models. If smaller residuals are needed, the accuracy of the DEIM model can be improved by either increasing k_{deim} or improving the accuracy of the reduced model. Increasing k_{deim} has negligible effect on online costs (since L^T is computed offline), but the benefits of doing this are limited. Figure 4.2 shows that the accuracy of DEIM solution is no better than the accuracy of the solution obtained using the (impractical) version of the reduced basis method given by

Algorithm 3; once that accuracy is obtained, increasing k_{deim} will produce no additional improvement. The accuracy of the reduced model can be improved using a stricter tolerance τ during the offline computation. Since choosing a stricter tolerance for the reduced model causes the size of the reduced basis k to increase, the cost of the reduced and DEIM models will also increase. This means that the benefits of more highly accurate DEIM computations will be obtained only for higher-resolution models.

In Table 4.1, the results for $n = 32$ and $m \geq 25$ are not shown. For these problems, the number of snapshots required to construct the reduced model, k_s exceeds the size of the pressure space $N_p = 3(n/2)^2 = 768$. This means that the number of snapshots required for the accuracy of the velocity is higher than the number of degrees of freedom in the full discretized pressure space. Therefore, the spatial discretization is not fine enough for reduced-order modeling to be necessary.

n_{fe}	m	4	9	16	25	36	49
32	k	306	942	1485			
	k_{deim}	102	314	495			
		time res	time res	time res	time res	time res	time res
	Full	1.11 <1.E-08	1.16 <1.E-08	1.05 <1.E-08			
	Reduced	0.44 3.78E-05	0.52 2.51E-05	0.66 8.89E-06			
	DEIM	0.15 2.22E-05	1.10 3.18E-05	3.40 4.61E-05			
64	k	273	825	1503	2394	3339	4455
	k_{deim}	91	275	501	798	1113	1485
		time res	time res	time res	time res	time res	time res
	Full	11.0 <1.E-08	10.8 <1.E-08	10.2 <1.E-08	11.3 <1.E-08	10.1 <1.E-08	10.3 <1.E-08
	Reduced	4.58 4.42E-05	5.07 2.99E-05	5.61 1.07E-05	5.75 1.06E-05	5.70 1.44E-05	5.42 3.97E-05
	DEIM	0.36 4.94E-05	2.38 1.67E-05	7.13 4.49E-05	19.4 4.90E-05	37.2 5.89E-05	70.8 7.42E-05
128	k	237	732	1383	2109	3039	4083
	k_{deim}	79	244	461	703	1013	1361
		time res	time res	time res	time res	time res	time res
	Full	135 <1.E-08	141 <1.E-08	147 <1.E-08	155 <1.E-08	132 <1.E-08	148 <1.E-08
	Reduced	55.1 3.24E-05	54.4 4.33E-05	53.6 4.50E-05	63.5 3.57E-05	66.8 3.35E-05	57.1 5.74E-05
	DEIM	1.14 1.38E-05	7.83 2.28E-05	22.2 2.76E-05	51.0 5.17E-05	101 5.33E-05	196 7.05E-05
		0.11 1.41E-05	0.51 2.47E-05	1.76 2.80E-05	4.84 5.27E-05	11.0 5.49E-05	24.8 7.16E-05

Table 4.1: Solution time and accuracy for Full, Reduced and DEIM models.

Figure 4.3 illustrates the tradeoff between accuracy and time for DEIM. The top plot compares the error indicators for an average $n_s = 10$ parameters and the bottom plot shows the CPU time for the two methods. While the cost of the DEIM does increase with the number of vectors k_{deim} , we reach similar accuracy as for the reduced model at a much lower cost. It can also be seen that the cost of increasing k_{deim} is small since the maximum considered here ($k_{deim} = 96$) is significantly smaller than $N = 1089$.

An alternative to the strategy of random sampling used to generate the reduced basis (Algorithm 5) is *greedy sampling*, which produces a basis of quasi-optimal dimension in the sense that when the Kolmogorov n -width exhibits algebraic or exponential decay in the error, the maximal error for greedy sampling differs from the Kolmogorov n -width by a constant factor. [5, 7]. Our previous experience [14], for parameter-dependent diffusion problems, is that the performance of the random sampling strategy used here is comparable to that of a greedy strategy. In particular, for several linear benchmark problems, to achieve comparable accuracy, we found that the size of the reduced basis generated by sampling was never more than 10% larger than that produced by a greedy algorithm and in many cases the basis sizes were identical. Moreover, the computational cost (in CPU time) of the random sampling strategy was significantly lower.

We explore this issue in Table 4.2 for the benchmark problem under consideration here. The table shows the size k_s of the snapshot set found by Algorithm 5 and the CPU times required to compute these basis vectors, together with the analogous results when a greedy algorithm is used with the same trial set and error indicator. These results were obtained for the benchmark problems for $m = 4$ parameters and three successively finer meshes, and, for one example, a smaller value of the tolerance τ for accepting the reduced basis. The following trends can be observed:

Fig. 4.3: Top: Error indicator for DEIM model versus k_{deim} . Bottom: CPU time vs. k_{deim} to solve using DEIM direct direct versus k_{deim} . For 32×32 finite element grid, $m = 4$ random variables, τ (tolerance for reduced basis) = 10^{-4} , k (rank of reduced basis) = 306, averaged over $n_s = 10$ samples.

- The random algorithm produces basis sets of larger dimension then the greedy algorithm but of similar order of magnitude; in three of the four cases, the dimension is larger by approximately 20% and in one ($n_{fe} = 128$) it is approximately 40%. These differences are somewhat larger than what we previously observed for diffusion problems.
- Random search is dramatically faster than greedy search, differing by at least a factor of 7 and at most 61.

The overhead for generating the reduced basis when there are more parameters (larger m) gets more significant. For example, the offline random search for the finest grid for $m = 49$ required approximately five days, and a greedy search would take dramatically longer. Our choice of random search represents a compromise between use of a slightly less-than-optimal basis and one that would be extremely expensive to compute. Our concern in this study is online strategies for reducing the cost of using the reduced model; it is clear that the online strategies considered here can be used for a reduced basis obtained using any method.

	$n_{fe} = 32, \tau = 10^{-4}$		$n_{fe} = 64, \tau = 10^{-4}$		$n_{fe} = 128, \tau = 10^{-4}$		$n_{fe} = 32, \tau = 10^{-6}$	
	k_s	CPU time	k_s	CPU time	k_s	CPU time	k_s	CPU time
Random	112	8.4min	96	35.5min	86	5.1hr	287	33.0min
Greedy	94	4.6hr	80	10.7hr	60	35.9 hr	233	33.6hr

Table 4.2: Sizes of snapshot sets (k_s) and CPU times for random and greedy generation of reduced basis.

4.2.1. Gappy POD. Another way to increase the accuracy of the reduced model is to increase the number of interpolation points in the approximation, while keeping the number of basis vectors fixed. This alternative to DEIM for selecting the indices is the so-called gappy POD method [8, 16, 34]. This method allows the number or rows selected by P^T to exceed the number of columns of V .

The approximation of the function using gappy POD looks similar to DEIM, but the inverse of $P^T V$ is replaced with the Moore-Penrose pseudoinverse $(P^T V)^\dagger$ [8]

$$(4.2) \quad \hat{F}(u) = V(P^T V)^\dagger P^T F(u).$$

To use this operator, we compute $P^T F(u)$ and solve the least squares problem

$$\alpha = \arg \min_{\hat{\alpha}} \|P^T V \hat{\alpha} - P^T F(u)\|_2 ,$$

which leads to the approximation $\hat{F}(u) = V\alpha$. Like $(P^T V)^{-1}$, the pseudoinverse can be precomputed, in this case using the SVD $P^T V = U\Sigma W^T$, giving

$$(P^T V)^\dagger = W\Sigma^\dagger U^T$$

where $\Sigma^\dagger$ is [22]

$$\begin{bmatrix} 1/\sigma_1 & & & \\ & \ddots & & \\ & & 1/\sigma_{k_{deim}} & \\ & & & 0 \end{bmatrix} .$$

With this approximation, the index selection method is described in Algorithm 6. Given V computed as in Algorithm 5, we require a method to determine the selection of the row indices that will lead to an accurate representation of the nonlinear component. The approach in [8] is an extension of the greedy algorithm used for DEIM (Algorithm 1). It takes as an input the number of grid points and the basis vectors. It simply chooses additional indices per basis vector where the indices correspond to the maximum index of the difference of the basis vector and its projection via the gappy POD model. Recall that in DEIM, the index associated with vector v_i that maximizes $|v_i - V(P^T V)^{-1} P^T v_i|$ (called ρ in Algorithm 1) is found, and P is augmented by e_ρ . In this extension of the algorithm, at each step i , this type of construction is done k_g/n_v times for v_i , where an index $\arg \max(|v_i - V(P^T V)^\dagger P^T v_i|)$ is found and then P and the projection of v_i are updated each time a new such index is chosen.

To compare the accuracy of this method with DEIM, we use Algorithm 5 to compute DEIM and modify line 21 to use Algorithm 6 with $k_g = 2k_{deim}$ for a range of values of k_{deim} , so that there are two indices selected for each v_i in the gappy algorithm. We use both methods to approximate the nonlinear component and solve the resulting models. Figure 4.4 shows the error indicators as functions of k_{deim} for both methods. It is evident that for smaller number of basis vectors the gappy POD provides additional accuracy. However, for larger number of basis vectors, the additional accuracy provided by gappy POD is small. Thus, the gappy POD method can be used to improve the accuracy when the number of basis vectors is limited. Since S is generated using the Full(k_s) described in Section 4.1, no additional accuracy is gained for the DEIM method when $k_{deim} > 102$.

5. Iterative methods. We have seen that the DEIM generates reduced-order models that produce solutions essentially as accurate as [those obtained by a version of reduced-order models that does not treat nonlinearities efficiently](#). In addition, Table 4.1 and Figure 4.3 illustrate that as expected, the DEIM model significantly decreases the online time spent constructing the nonlinear component of the reduced model. Since $Q_u^T C(\vec{u}) Q_u$ has been replaced by a cheap approximation, $L^T P^T C(\vec{u}) Q_u$, the remaining cost of the nonlinear iteration in the DEIM is that of the linear system solve in line 2 of Algorithm 4.

The cost of this computation depends on the rank of the reduced basis k , which is $3k_s$ in our setting, where k_s is the number of snapshots used to construct the reduced basis. These quantities depend on properties of the problem such as the number of parameters in the model or the desired level of accuracy in the reduced model. In contrast, the cost of solving the full model is independent of the number of parameters, and it could be as small as $O(N)$ if multigrid methods can be utilized. Thus, it may happen that k is much less than N , but the $O(k^3)$ cost of solving the DEIM model is larger than the cost of solving the full model. This is the case, for example, for $n_{fe} = 64$, $m = 49$ in Table 4.1, where the CPU time to solve the full model is half that of solving the DEIM model using direct methods.

An alternative is to use iterative methods to solve the reduced model. Their cost is $O(k^2 p)$ where p is the number of iterations required for convergence, so that there are values of k where, if p is small enough, iterative methods will be preferable to direct methods. We have seen examples of this for linear problems in [14]. In this section, we discuss the use of iterative methods based on preconditioned Krylov subspace methods to improve the efficiency of the DEIM model.

Algorithm 6 Index selection using gappy POD [8]

Input: k_g number of indices to choose, $V = [v_1, \dots, v_{n_v}]$, an $N \times n_v$ matrix with columns made up of the left singular vectors from the POD of the nonlinear snapshot matrix S .

Output: P , extracts the indices used for the interpolation.

```

1:  $n_b = 1$ ,  $n_{it} = \min(n_v, k_g)$   $n_{c,\min} = \left\lfloor \frac{n_v}{n_{it}} \right\rfloor$   $n_{a,\min} = \left\lfloor \frac{k_g}{n_v} \right\rfloor$ 
2: for  $i = 1, \dots, n_{it}$  do
3:    $n_c = n_{c,\min}$ ,  $n_a = n_{a,\min}$ 
4:   if  $i \leq (n_v \bmod n_{it})$  then  $n_c = n_c + 1$  end if
5:   if  $i \leq (k_g \bmod n_v)$  then  $n_a = n_a + 1$  end if
6:   if  $i == 1$  then
7:      $r = \sum_{q=1}^{n_c} v_q^2$ 
8:     for  $j = 1, \dots, n_a$  do  $\rho_j = \text{argmax}(r)$ ,  $r[\rho_j] = 0$ 
9:   end for
10:   $P = [e_{\rho_1}, \dots, e_{\rho_{n_a}}]$ ,  $\hat{V} = [v_1, \dots, v_{n_c}]$ 
11: else
12:   for  $q = 1, \dots, n_c$  do
13:      $\alpha = \min_{\hat{\alpha}} \|P^T \hat{V} \hat{\alpha} - P^T v_{n_b+q}\|_2$ 
14:      $R_q = v_{n_b+q} - \hat{V} \alpha$ 
15:   end for
16:    $r = \sum_{q=1}^{n_c} R_q^2$ 
17:   for  $j = 1, \dots, n_a$  do
18:      $\rho_j = \text{argmax}(r)$ ,  $P = [P, e_{\rho_j}]$ 
19:     for  $q = 1, \dots, n_c$  do
20:        $\alpha = \min_{\hat{\alpha}} \|P^T \hat{V} \hat{\alpha} - P^T v_{n_b+q}\|_2$ 
21:        $R_q = v_{n_b+q} - \hat{V} \alpha$ 
22:     end for
23:      $r = \sum_{q=1}^{n_c} R_q^2$ 
24:   end for
25:    $\hat{V} = [\hat{V}, v_{n_b+1}, \dots, v_{n_b+n_c}]$ ,  $n_b = n_b + n_c$ 
26: end if
27: end for

```

For iterative methods to be efficient, effective preconditioners are needed. In the offline-online paradigm, it is also desirable to make the construction of the preconditioner independent of parameters, so that this construction can be part of the offline step. Thus, we will develop preconditioners that depend only on the mean value $\xi^{(0)}$ from the parameter set, and refer to such techniques as “offline” preconditioners. To see the impact of this choice, we will also compare the performance of these approaches with “online” versions of them, where the preconditioning operator for a model with parameter ξ is built using that parameter; this approach is not meant to be used in practice since its online cost will depend on N , but it provides insight concerning a lower bound on the iteration count that can be achieved using an offline preconditioner. Versions of the offline approach have been used with stochastic Galerkin methods in [25].

We consider two preconditioners of the DEIM model:

1. *The Stokes preconditioner* is the matrix used for the reduced Stokes solve in equation (3.6),

$$(5.1) \quad M_r(\xi) = Q^T \begin{bmatrix} A(\xi) & B^T \\ B & 0 \end{bmatrix} Q.$$

2. *The Navier-Stokes preconditioner* uses the converged solution of the full model $\vec{u}_n$ as the input for $P^T C(\vec{u}_n)$ and uses the operator from the DEIM model from equation (3.7),

$$(5.2) \quad M_r(\xi) = Q^T \begin{bmatrix} A(\xi) & B^T \\ B & 0 \end{bmatrix} Q + \begin{bmatrix} L^T P^T C(\vec{u}_n) Q_u & 0 \\ 0 & 0 \end{bmatrix}.$$

Fig. 4.4: Average error indicator as a function of basis vectors for reduced, DEIM, and gappy POD methods. For 32×32 finite element grid, $m = 4$ random variables, τ (tolerance for reduced basis) = 10^{-4} , k (rank of reduced basis) = 306, varying k_{deim} and $k_g = 2k_{deim}$. Averaged over $n_s = 100$ samples.

In both cases, the preconditioned coefficient matrix is $\mathcal{A}(\xi)M_r^{-1}$ where $\mathcal{A}(\xi)$ is the coefficient matrix in (3.7). According to the comments above, we consider two versions of each of these, the offline version, $M_r(\xi_0)$, and the online version, $M_r(\xi)$. Clearly, for the first step of the reduced nonlinear iteration, a Stokes solve, only one linear iteration will be required using the online Stokes preconditioner.

For these experiments, we solve the steady-state Navier-Stokes equations for the driven cavity flow problem using the full model, reduced model, and the DEIM model. For the DEIM model, the linear systems are solved using both direct and iterative methods. The iterative methods use the preconditioned **bicgstab** method [32].

The offline construction is described in Algorithm 5; we use $\tau = 10^{-4}$ and $k_{trial} = 2000$. The algorithm chooses k_s snapshots and produces Q_u of rank $2k_s$ and Q_p of rank k_s , yielding reduced models of rank $k = 3k_s$. The online experiments are run for $n_s = 10$ random parameters. The average number of iterations required for the convergence of the linear systems is presented in Table 5.1 and the average time for the entire nonlinear solve of each model is presented in Table 5.2. The nonlinear solve time includes the time to compute C or $P^T C$, but not the time for assembly of the linear component of the model nor the computation time for η_ξ , the norm of the full residual (3.4) for the approximate solution found by the reduced and DEIM models. The iterative methods presented in this table use offline preconditioners. The lowest online CPU time is in bold. The nonlinear iterations are run to tolerance $\delta = 10^{-8}$ and the **bicgstab** method for a reduced system with coefficient matrix, $Q^T \mathcal{A} Q$, stops when the solution $x^{(i)}$ satisfies

$$\frac{\|r - Q^T \mathcal{A} Q x^{(i)}\|}{\|r\|} < 10^{-9}$$

where r is $Q^T \begin{bmatrix} \mathbf{f} \\ 0 \end{bmatrix}$.

Table 5.1 illustrates that the performance of the offline preconditioners using the mean parameter compared to the versions that use the exact parameter, with the offline preconditioners requiring more iterations, as expected. In Table 5.2 we compare these offline parameters with the direct DEIM method and determine that for large enough m (number of parameters) and k (size of the reduced basis), the iterative methods

n_{fe}	m (# parameters)		4	9	16	25	36	49
32	k (reduced basis rank)		306	942	1485			
	k_{deim} (# DEIM vectors)		102	314	495			
	Offline Stokes		11.3	16.2	19.8			
	Online Stokes		2.0	2.3	2.3			
	Offline Navier-Stokes		11.5	16.1	20.5			
	Online Navier-Stokes		1.8	1.9	1.9			
64	k		273	825	1503	2394	3339	4455
	k_{deim}		91	275	501	798	1113	1485
	Offline Stokes		10.3	13.9	16.9	17.6	19.9	23.3
	Online Stokes		2.1	2.1	2.3	2.4	2.2	2.4
	Offline Navier-Stokes		10.5	13.6	16.5	17.3	19.9	24.0
	Online Navier-Stokes		1.7	1.8	1.9	2.0	1.9	2.0
128	k		237	732	1383	2109	3039	4083
	k_{deim}		79	244	461	703	1013	1361
	Offline Stokes		8.8	16.4	21.0	17.8	19.9	25.4
	Online Stokes		1.9	2.2	2.5	2.3	2.2	2.4
	Offline Navier-Stokes		8.9	16.5	20.5	17.9	20.1	25.1
	Online Navier-Stokes		1.8	1.8	2.1	2.1	1.9	2.1

Table 5.1: Average iteration counts of preconditioned **bicgstab** for solving equation (3.7), over $n_s = 10$ parameter samples.

n_{fe}	m (# parameters)		4	9	16	25	36	49
32	k (reduced basis rank)		306	942	1485			
	k_{deim} (# DEIM vectors)		102	314	495			
	Full	Direct	1.11	1.16	1.06			
	Reduced	Direct	0.18	1.21	3.00			
	DEIM	Direct	0.05	0.37	1.02			
	DEIM	Stokes	0.05	0.42	0.95			
64	DEIM		0.05	0.44	1.03			
	k		273	825	1503	2394	3339	4455
	k_{deim}		91	275	501	798	1113	1485
	Full	Direct	11.0	10.8	10.2	11.3	10.1	10.3
	Reduced	Direct	0.48	2.53	7.33	20.5	39.5	76.2
	DEIM	Direct	0.07	0.27	1.08	4.40	9.54	20.7
128	DEIM	Stokes	0.07	0.29	0.86	2.29	4.62	9.04
	DEIM	Navier-Stokes	0.08	0.30	0.87	2.27	4.39	8.98
	k		237	732	1383	2109	3039	4083
	k_{deim}		79	244	461	703	1013	1361
	Full	Direct	135	141	147	155	132	148
	Reduced	Direct	1.62	7.25	23.8	56.7	98.1	191
256	DEIM	Direct	0.09	0.39	1.12	3.59	7.11	15.7
	DEIM	Stokes	0.09	0.45	1.10	2.40	3.89	8.74
	DEIM	Navier-Stokes	0.09	0.45	1.08	2.45	3.91	8.62

n_{fe}	m (# parameters)	4	9	16	25	36	49
32	k (reduced basis rank)	306	942	1485			
	k_{deim} (# DEIM vectors)	102	314	495			
	Stokes	0.04	0.36	1.02			
	Navier-Stokes	1.01	1.42	2.06			
64	k	273	825	1503	2394	3339	4455
	k_{deim}	91	275	501	798	1113	1485
	Stokes	0.10	0.85	2.38	6.51	13.7	26.7
	Navier-Stokes	8.96	10.0	12.6	15.7	22.8	35.7
128	k	237	732	1383	2109	3039	4083
	k_{deim}	79	244	461	703	1013	1361
	Stokes	0.29	1.96	6.84	17.9	33.0	68.2
	Navier-Stokes	116	132	129	143	147	193

Table 5.3: CPU time to construct the preconditioner

are faster than direct methods. For example, for $m = 9$, the direct methods are slightly faster whereas for $m = 16$ the iterative methods are faster for all values of n_{fe} . We also note that the fastest DEIM method is faster than the full model for all cases. Returning to the motivating example of $n = 64$ and $m = 49$, the DEIM iterative method is faster than the full model, whereas the DEIM direct method performs twice as slowly as the full model. Thus, utilizing iterative methods increases the range of k where reduced-order modeling is practical.

Table 5.3 presents the (offline) cost of constructing the preconditioners. Since the Navier-Stokes preconditioner uses the full solution, the cost of constructing this preconditioner scales with the costs of the full solution. The cost of constructing the Stokes preconditioner is significantly smaller, because it does not require solving the full nonlinear problem at the mean parameter. It performs similarly to the exact preconditioner in the online computations. Thus, the exact Stokes preconditioner is an efficient option for both offline and online components of this problem.

Remark. A variant of the offline preconditioning methods discussed here is a *blended* approach, in which a small number of preconditioners corresponding to a small set of parameters is constructed offline. Then for online solution of a problem with parameter ξ , the preconditioner derived from the parameter closest to ξ can be applied. We found that this approach did not improve performance of the preconditioners considered here.

6. Conclusion. We have shown that the discrete interpolation method is effective for solving the steady-state Navier-Stokes equations. This approach produces a reduced-order model that is essentially as accurate as a naive implementation of a reduced basis method without incurring online costs of order N . In cases where the dimension of the reduced basis is large, performance of the DEIM is improved through the use of preconditioned iterative methods to solve the linear systems arising at each nonlinear Picard iteration. This is achieved using the mean parameter to construct preconditioners. These preconditioners are effective for preconditioning the reduced model in the entire parameter space.

Acknowledgement. We thank three anonymous referees for careful reading of the manuscript, which led to improvements in the presentation.

REFERENCES

- [1] D. Amsallem, M. Zahr, Y. Choi, and C. Farhat. Design optimization using hyper-reduced-order models. *Structural and Multidisciplinary Optimization*, 51:919–940, 2015.
- [2] H. Antil, M. Heinkenschloss, and D. C. Sorensen. Application of the discrete empirical interpolation method to reduced order modeling of nonlinear and parametric systems. In G. Rozza, editor, *Springer MSEA series: Reduced Order Methods for Modeling and Computational Reduction*, volume 8. Springer-Verlag, Italia, Milano, 2013.
- [3] F. Ballarin, A. Manzoni, A. Quarteroni, and G. Rozza. Supremizer stabilization of POD-Galerkin approximation of

- parametrized steady incompressible Navier-Stokes equations. *International Journal for Numerical Methods in Engineering*, 102:1136–1161, 2015.
- [4] M. Barrault, Y. Maday, N. C. Nguyen, and A. T. Patera. An ‘empirical interpolation’ method: application to efficient reduced-basis discretization of partial differential equations. *Comptes Rendus Mathématique*, 339(9):667–672, 2004.
 - [5] P. Binev, A. Cohen, W. Dahmen, R. DeVore, G. Petrova, and P. Wojtaszczyk. Convergence rates for greedy algorithms in reduced basis methods. *SIAM Journal on Mathematical Analysis*, 43:1457–1472, 2011.
 - [6] P. Bochev and R. B. Lehoucq. On the finite element solution of the pure Neumann problem. *SIAM Review*, 47(1):50–66, 2005.
 - [7] A. Buffa, Y. Maday, A. T. Patera, C. Prud’homme, and G. Turinici. A priori convergence of the greedy algorithm for the parametrized reduced basis method. *ESAIM: Mathematical Modelling and Numerical Analysis*, 46(03):595–603, 2012.
 - [8] K. Carlberg, C. Farhat, J. Cortial, and D. Amsallem. The GNAT method for nonlinear model reduction: Effective implementation and application to computational fluid dynamics and turbulent flows. *Journal of Computational Physics*, 242:623–647, 2013.
 - [9] S. Chaturantabut and D. C. Sorensen. Nonlinear model reduction via discrete empirical interpolation. *SIAM Journal on Scientific Computing*, 32(5):2737–2764, 2010.
 - [10] T. A. Davis. Algorithm 832: UMFPACK v4. 3—an unsymmetric-pattern multifrontal method. *ACM Transactions on Mathematical Software (TOMS)*, 30(2):196–199, 2004.
 - [11] S. Deparis. Reduced basis error bound computation of parameter-dependent Navier-Stokes equations by the natural norm approach. *SIAM Journal on Numerical Analysis*, 46:2039–2067, 2008.
 - [12] S. Deparis and A. E. Løvgrén. Stabilized reduced basis approximation of incompressible Navier-Stokes equations. *Journal of Scientific Computing*, 50:198–212, 2012.
 - [13] S. Deparis and G. Rozza. Reduced basis method for multi-parameter-dependent steady Navier-Stokes equations: Applications to natural convection in a cavity. *Journal of Computational Physics*, 228:4359–4378, 2009.
 - [14] H. C. Elman and V. Forstall. Preconditioning techniques for reduced basis methods for parameterized elliptic partial differential equations. *SIAM Journal on Scientific Computing*, 37(5):S177–S194, 2015.
 - [15] H. C. Elman, D. Silvester, and A. Wathen. *Finite Elements and Fast Iterative Solvers: with Applications in Incompressible Fluid Dynamics*. Oxford University Press, 2014.
 - [16] R. Everson and L. Sirovich. Karhunen–Loève procedure for gappy data. *Journal of the Optical Society of America A*, 12(8):1657–1664, 1995.
 - [17] V. Girault and P. A. Raviart. *Finite Element Approximation of the Navier-Stokes Equations*. Springer-Verlag, New York, 1986.
 - [18] M. A. Grepl, Y. Maday, N. C. Nguyen, and A. T. Patera. Efficient reduced-basis treatment of nonaffine and nonlinear partial differential equations. *ESAIM: Mathematical Modelling and Numerical Analysis*, 41(03):575–605, 2007.
 - [19] J. Kim. Phase field computations for ternary fluid flows. *Computer Methods in Applied Mechanics and Engineering*, 196(45):4779–4788, 2007.
 - [20] Y. Maday, N. C. Nguyen, A. T. Patera, and G. S. H. Pau. A general multipurpose interpolation procedure: the magic points. *Communications on Pure and Applied Analysis*, 8:383–404, 2009.
 - [21] A. Manzoni. An efficient computational framework for reduced basis approximation and a posteriori error estimation of parametrized Navier-Stokes flows. *ESAIM: Mathematical Modelling and Numerical Analysis*, 48:1199–1226, 2014.
 - [22] D. P. O’Leary. *Scientific Computing with Case Studies*. SIAM, 2009.
 - [23] M. A. Olshanskii and A. Reusken. Analysis of a Stokes interface problem. *Numerische Mathematik*, 103(1):129–149, 2006.
 - [24] J. S. Peterson. The reduced basis method for incompressible viscous flow calculations. *SIAM J. Sci. Stat. Comput.*, 10:777–786, 1989.
 - [25] C. E. Powell and D. J. Silvester. Preconditioning steady-state Navier–Stokes equations with random data. *SIAM Journal on Scientific Computing*, 34(5):A2482–A2506, 2012.
 - [26] A. Quarteroni, A. Manzoni, and F. Negri. *Reduced Basis Methods for Partial Differential Equations, An Introduction*. Springer, 2016.
 - [27] A. Quarteroni and G. Rozza. Numerical solution of parametrized Navier-Stokes equations by reduced basis methods. *Numerical Methods for Partial Differential Equations*, 23(4):923–948, 2007.
 - [28] G. Veroy and K. Veroy. On the stability of the reduced basis method for Stokes equations in parametrized domains. *Computer Methods in Applied Mechanics and Engineering*, 196(7):1244–1260, 2007.
 - [29] D. Silvester. IFISS 3.3 release notes. <http://www.maths.manchester.ac.uk/%7Edjs/ifiss/release.txt>, 2013.
 - [30] D. Silvester, H. C. Elman, and A. Ramage. Incompressible Flow and Iterative Solver Software (IFISS) version 3.4, July 2015. <http://www.manchester.ac.uk/ifiss/>.
 - [31] Z. Tan, D. V. Le, Z. Li, K. M. Lim, and B. C. Khoo. An immersed interface method for solving incompressible viscous flows with piecewise constant viscosity across a moving elastic membrane. *Journal of Computational Physics*, 227(23):9955–9983, 2008.
 - [32] H. A. van der Vorst. *Iterative Krylov Methods for Large Linear Systems*. Cambridge University Press, New York, 2003.
 - [33] K. Veroy and A. T. Patera. Certified real-time solution of the parametrized steady incompressible Navier-Stokes equations: rigorous reduced-basis a posteriori error bounds. *International Journal for Numerical Methods in Fluids*, 47(8):773–788, 2005.
 - [34] K. Willcox. Unsteady flow sensing and estimation via the gappy proper orthogonal decomposition. *Computers and Fluids*, 35:208226, 2006.
 - [35] D. Xiao, F. Fang, A.G. Buchan, C. C. Pain, I. M. Navon, J. Du, and G. Hu. Non-linear model reduction for the Navier-Stokes equations using the residual DEIM method. *Journal of Computational Physics*, 263:1–18, 2014.