



Application of adaptive ANOVA and reduced basis methods to the stochastic Stokes-Brinkman problem

Kevin Williamson¹  · Heyrim Cho² · Bedřich Sousedík¹

Received: 27 August 2020 / Accepted: 15 February 2021 / Published online: 17 March 2021
© The Author(s), under exclusive licence to Springer Nature Switzerland AG part of Springer Nature 2021

Abstract

The Stokes-Brinkman equations model fluid flow in highly heterogeneous porous media. In this paper, we consider the numerical solution of the Stokes-Brinkman equations with stochastic permeabilities, where the permeabilities in subdomains are assumed to be independent and uniformly distributed within a known interval. We employ a truncated anchored ANOVA decomposition alongside stochastic collocation to estimate the moments of the velocity and pressure solutions. Through an adaptive procedure selecting only the most important ANOVA directions, we reduce the number of collocation points needed for accurate estimation of the statistical moments. However, for even modest stochastic dimensions, the number of collocation points remains too large to perform high-fidelity solves at each point. We use reduced basis methods to alleviate the computational burden by approximating the expensive high-fidelity solves with inexpensive approximate solutions on a low-dimensional space. We furthermore develop and analyze rigorous a posteriori error estimates for the reduced basis approximation. We apply these methods to 2D problems considering both isotropic and anisotropic permeabilities.

Keywords Stokes-Brinkman equations · Reduced basis methods · ANOVA · Porous media · A posteriori error estimates

Mathematics Subject Classification (2010) 35R60 · 65C30 · 60H35 · 65N75

1 Introduction

The simulation of flow in porous media has numerous applications, to include reservoir simulation, nuclear waste disposal, and carbon dioxide sequestration. Such simulation is challenging for a variety of reasons. First, the domains tend to be fairly irregular, which complicates the model

Bedřich Sousedík was partially supported by the U.S. National Science Foundation under grant DMS 1913201.

✉ Kevin Williamson
kwillia1@umbc.edu

Heyrim Cho
heyrim.cho@ucr.edu

Bedřich Sousedík
sousedik@umbc.edu

¹ Department of Mathematics and Statistics, University of Maryland, Baltimore County, 1000 Hilltop Circle, Baltimore, MD 21250, USA

² Department of Mathematics, University of California Riverside, Riverside, CA 92521, USA

geometry. Second, the geologic formations consist of many varying materials, with different geologic properties. Third, there are often fractures and vugs within the domain that alter the effective permeabilities. The standard approach to modeling these types of problems is to couple Darcy's law and the Stokes equations and enforce the Beavers-Joseph-Saffman conditions along the interface [1–3]. The free-flow regions (fractures, vugs) are modeled using Stokes flow, whereas the porous region is modeled using Darcy's law [4, 5]. However, the two types of domains are not well-separated in reservoirs, and it may be difficult to determine the appropriate conditions to enforce along the interface. We model flow in porous media using the Stokes-Brinkman equations [6–10], which combine the Stokes equations and Darcy's law into a single system of equations. The Stokes-Brinkman equations reduce to Stokes or Darcy flow depending upon the coefficients and were suggested as a replacement for the coupled Stokes-Darcy equations in [10]. By careful selection of coefficients, the equations allow modeling of free-flow and porous domains together, thereby resolving issues along the interface.

In this paper, we consider the case where the exact permeabilities are unknown but are instead specified by

a known probability distribution. The first step in many methods to solve such stochastic PDEs is to parametrize the distribution by a finite number of parameters. In some instances, such as the case studied in this paper, such a parametrization arises naturally from a partitioning of the domain into a finite number of subdomains. In other instances, a truncated Karhunen-Loëve expansion [11, 12] may be used to obtain such a finite parametrization. Once the stochastic space has been parametrized, statistical moments may be approximated through stochastic collocation [13] in which deterministic solves are performed at carefully selected points.

The straightforward stochastic collocation method suffers from the curse of dimensionality—complexity grows exponentially with the number of parameters. To reduce the computational complexity of collocation methods, one may utilize Smolyak sparse grid [14] or an ANOVA decomposition [15–17]. We consider the latter in this paper. By using a truncated ANOVA decomposition, we may resolve the curse of dimensionality by decomposing the intractable high-dimensional problem into a set of tractable problems of low stochastic dimension which we solve by stochastic collocation. We obtain further improvements by adaptively selecting the ANOVA terms [18, 19].

Whereas a truncated ANOVA decomposition or Smolyak sparse grid significantly reduces the number of collocation points, the number of remaining collocation points might still be computationally prohibitive if the cost of solving at a single collocation point is high. Such is the case with the discrete Stokes-Brinkman systems, which are highly ill-conditioned due to the heterogeneity in permeabilities. Reduced basis methods [20–23] alleviate this computational burden by approximating the manifold of solutions by a low-dimensional linear space. An expensive offline step is first performed to generate the reduced basis. However, after the reduced basis is built, all subsequent solves at collocation points are replaced by cheap low-dimensional computations on the reduced basis.

Reduced basis methods in conjunction with ANOVA for parametric partial differential equations were first studied in [24] where a three step RB-ANOVA-RB method was suggested for solving high-dimensional problems. The first two steps (RB-ANOVA) are similar to the idea presented in this paper in which a reduced basis is built and then used to compute the ANOVA expansion. However, in our work, the construction of the reduced basis and the ANOVA expansion occur simultaneously rather than as separate steps. Furthermore, in [24], only a first-level ANOVA expansion is considered for the purpose of identifying parameters which have low sensitivity and can be given fixed values, thereby reducing the overall parametric dimension for the reduced basis solve in the final (RB) step. Instead, we adaptively select the ANOVA terms to allow higher-level

terms at minimal cost. Key to the success of our method is the use of anchored ANOVA [18, 19, 25–27] which allows us to replace the high-dimensional integrals of the ANOVA expansion with cheaper low-dimensional integrals. In this way, our work is more closely aligned with the work in [28, 29] where reduced basis methods are used in conjunction with anchored ANOVA and adaptive selection of ANOVA terms to solve high-dimensional stochastic partial differential equations.

In this paper, we apply these methods to the stochastic Stokes-Brinkman problem. As a saddle-point problem, the Stokes-Brinkman equations introduce additional complexity to the reduced basis methods because the reduced basis systems must be inf-sup stable. To guarantee inf-sup stability, we follow the approach of [30], which considers the simpler Stokes case. We also devise rigorous a posteriori error estimates based upon the Brezzi stability theory, following the approach of [31]. These a posteriori error estimates enable us to be confident in the accuracy of the reduced basis approximations and are useful in building the reduced basis.

This paper is organized as follows. In Section 2, we introduce the Stokes-Brinkman equations, present their discretization using mixed finite elements, and describe the parametrization of the stochastic permeabilities. In Section 3, we discuss the ANOVA decomposition and the adaptive selection of terms. An overview of reduced basis methods is presented in Section 4, and the rigorous a posteriori error estimates upon which these methods rely are presented in Section 5. Finally, in Section 6, we present numerical experiments demonstrating the effectiveness of these techniques for stochastic Stokes-Brinkman problems, considering problems with both isotropic and anisotropic permeabilities.

2 The Stokes-Brinkman problem and its discretization

2.1 The Stokes-Brinkman equations

Let $D \subset \mathbb{R}^d$, $d = 2, 3$, be a connected, open domain with Lipschitz boundary ∂D . The Stokes-Brinkman equations [6–10] model the flow of a viscous fluid in heterogeneous porous material as

$$-\nu^* \Delta u + \nu \mathbf{K}^{-1} u + \nabla p = f \quad (1)$$

$$\nabla \cdot u = 0, \quad (2)$$

where $\nu > 0$ is the constant viscosity of the fluid, $\nu^* > 0$ is an effective viscosity, \mathbf{K} is a symmetric positive definite permeability tensor, u is the velocity, p is the pressure, and f denotes external forces. In this paper, we restrict

attention to the case where \mathbf{K} is a diagonal matrix. We further require that v^* and \mathbf{K}^{-1} are bounded above on D . Equation 1 is derived from the conservation of momentum, and Eq. 2 is derived from the conservation of mass. They are accompanied by the following Dirichlet and Neumann boundary conditions

$$\begin{aligned} u &= u_D & \text{on } \partial D_D \\ v^* \frac{\partial u}{\partial n} - p n &= u_N & \text{on } \partial D_N, \end{aligned}$$

where ∂D_D and ∂D_N denote the Dirichlet and Neumann boundaries, respectively, n is the unit outward normal, and $\frac{\partial u}{\partial n}$ is the directional derivative of the velocity in the normal direction.

The Stokes-Brinkman equations may be understood as limiting cases of the Stokes equations and Darcy's law [4, 5]. Indeed, if $\mathbf{K} \gg 0$ and $v^* = v$, Eq. 1 approximates the Stokes equation $-v \Delta u + \nabla p = f$, and, as $v^* \rightarrow 0$, it approximates Darcy's law $u = -\frac{\mathbf{K}}{v}(\nabla p - f)$. Thus, the Stokes-Brinkman equations provide a single system of equations to solve in highly heterogeneous porous media. In general, the choice of v^* will depend upon the porosity of the material. However, for small permeabilities, the diffusive term $v^* \Delta u$ introduces only a small perturbation to Darcy's law [9]. In the absence of precise information about the value of v^* in regions of small permeability, it is common to choose $v^* = v$, which is the convention we adopt in this paper. For more details on the choice of v^* , see [9, pp. 26–29].

2.2 Finite element discretization

We seek a weak solution to Eqs. 1 and 2. Let $L^2(D)$ denote the space of square-integrable functions on D ,

$$L^2(D) \equiv \{q: \int_D q^2 < \infty\},$$

and $H^1(D)$ denote the subspace of $L^2(D)$ with weak derivatives in $L^2(D)$,

$$H^1(D) \equiv \{q \in L^2(D): \frac{\partial q}{\partial x_i} \in L^2(D), i = 1, \dots, d\}.$$

The space $H^1(D)^d$ denotes the space of vector-valued functions whose d components are each in $H^1(D)$. We define the following spaces for the velocity

$$\begin{aligned} H_E^1(D) &\equiv \{u \in H^1(D)^d: u = u_D \text{ on } \partial D_D\} \\ H_{E_0}^1(D) &\equiv \{u \in H^1(D)^d: u = 0 \text{ on } \partial D_D\}. \end{aligned}$$

We choose a fixed $w \in H_E^1(D)$ and note that any $v \in H_E^1(D)$ may be written uniquely as $u + w$ for some $u \in H_{E_0}^1(D)$.

We define the bilinear forms

$$a_S(u, v) \equiv \int_D v^* \nabla u : \nabla v, \quad u, v \in H^1(D)^d \quad (3)$$

$$a_D(u, v) \equiv \int_D u \cdot (\mathbf{K}^{-1} v), \quad u, v \in H^1(D)^d \quad (4)$$

$$a(u, v) \equiv a_S(u, v) + a_D(u, v), \quad u, v \in H^1(D)^d \quad (5)$$

$$b(v, q) \equiv - \int_D q \nabla \cdot v, \quad v \in H^1(D)^d, q \in L^2(D) \quad (6)$$

where $\nabla u : \nabla v = \sum_{i=1}^d \nabla u_i \cdot \nabla v_i$. The bilinear forms $a_S(\cdot, \cdot)$ and $a_D(\cdot, \cdot)$ denote the Stokes and Darcy parts, respectively, of the bilinear form $a(\cdot, \cdot)$. We define the linear functionals

$$\ell_1(v) \equiv \int_D f \cdot v - a(w, v) + \int_{\partial D_N} u_N \cdot v, \quad v \in H_{E_0}^1(D), \quad (7)$$

$$\ell_2(q) \equiv -b(w, q), \quad q \in L^2(D). \quad (8)$$

The weak formulation of Eqs. 1–2 is then to find $u \in H_{E_0}^1(D)$ and $p \in L^2(D)$ such that

$$a(u, v) + b(v, p) = \ell_1(v), \quad \forall v \in H_{E_0}^1(D) \quad (9)$$

$$b(u, q) = \ell_2(q), \quad \forall q \in L^2(D). \quad (10)$$

The full velocity solution with proper Dirichlet boundary conditions is then $u + w$.

In the following, we denote the space $H_{E_0}^1(D)$ by V and the space $L^2(D)$ by Q . Furthermore, due to the presence of zero Dirichlet boundary conditions, we define the V inner product as $(u, v)_V \equiv \int_D \nabla u : \nabla v$.

The existence and uniqueness of solutions to Eqs. 9 and 10 is a consequence of the Brezzi stability conditions [32, 33]. We state these conditions in Theorem 1, paraphrasing Corollary 4.2.1 of [32] for the case of a symmetric $a(\cdot, \cdot)$. This theorem additionally provides stability estimates which we will use to derive our a posteriori error estimates in Section 5.

Theorem 1 *Let V and Q be Hilbert Spaces. Let $a: V \times V \rightarrow \mathbb{R}$ be a continuous symmetric bilinear form that satisfies the coercivity condition: there exists an $\alpha > 0$ such that $a(u, u) \geq \alpha \|u\|_V$ for all $u \in V$. Furthermore, let $b: V \times Q \rightarrow \mathbb{R}$ be a continuous bilinear form that satisfies the inf-sup condition: there exists $\beta > 0$ such that*

$$\inf_{q \in Q} \sup_{v \in V} \frac{b(v, q)}{\|q\|_Q \|v\|_V} \geq \beta. \quad (11)$$

Then, for any $\ell_1 \in V^$ and $\ell_2 \in Q^*$, the saddle-point system (9)–(10) has a unique solution $(u, p) \in V \times Q$ which satisfies the stability bounds*

$$\|u\|_V \leq \frac{1}{\alpha} \|\ell_1\|_{V^*} + \frac{2}{\beta} \sqrt{\frac{\gamma}{\alpha}} \|\ell_2\|_{Q^*} \quad (12)$$

$$\|p\|_Q \leq \frac{2}{\beta} \sqrt{\frac{\gamma}{\alpha}} \|\ell_1\|_{V^*} + \frac{\gamma}{\beta^2} \|\ell_2\|_{Q^*} \quad (13)$$

where α and β are the above coercivity and inf-sup constants and $\gamma > 0$ is the continuity constant satisfying $a(u, v) \leq \gamma \|u\|_V \|v\|_V$ for all $u, v \in V$.

Coercivity and continuity of Eq. 3 and continuity of Eq. 6 are immediately seen to be satisfied as these are the forms used in the Stokes equations. The fact that \mathbf{K}^{-1} is symmetric positive definite implies that $a(u, u) = a_S(u, u) + a_D(u, u) > a_S(u, u)$ and coercivity of Eq. 5 immediately follows. Continuity of Eq. 5 follows as the bilinear forms $a_S(\cdot, \cdot)$ and $a_D(\cdot, \cdot)$ are both bounded due to boundedness of the viscosities and inverse permeabilities. The inf-sup condition is more delicate, but the spaces V and Q above were chosen to satisfy this property [34, Chapter 3].

Remark 1 In the case where no Neumann boundary conditions are specified, the inf-sup condition (11) fails as $b(q, v) = 0$ for all $v \in V$ where q is constant. In this case, we may define Q to be the quotient space $L^2(D)/\mathbb{R}$ in which two functions $q_1, q_2 \in L^2(D)$ are identified if they differ by a constant function. By equipping this space with the norm $\|q\|_Q = \|q - 1/|D| \int_D q\|_{L^2(D)}$, the inf-sup condition is satisfied. This demonstrates that the pressure solution in this case is unique up to a constant. In this paper, we will always specify Neumann conditions on a portion of the boundary.

We solve (9) and (10) using the mixed finite element method [32]. Given conforming finite element spaces $V_h \subset V$ and $Q_h \subset Q$, we seek $(u_h, p_h) \in V_h \times Q_h$ such that

$$\begin{aligned} a(u_h, v_h) + b(v_h, p_h) &= \ell_1(v_h), \quad \forall v_h \in V_h, \\ b(u_h, q_h) &= \ell_2(q_h), \quad \forall q_h \in Q_h. \end{aligned}$$

By selecting finite element bases for V_h and Q_h , we obtain the discrete saddle-point problem

$$\begin{bmatrix} \mathbf{A} & \mathbf{B}^T \\ \mathbf{B} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} \mathbf{f} \\ \mathbf{g} \end{bmatrix}, \quad (14)$$

where \mathbf{A} and \mathbf{B} are discrete analogs of Eqs. 5 and 6, respectively, \mathbf{u} and \mathbf{p} are the vectors of degrees of freedom for the velocity and pressure, respectively, and \mathbf{f} and \mathbf{g} are the discretizations of the functionals (7) and (8), respectively.

The coercivity and continuity properties are immediately inherited by the finite element discretization. However, care must be taken to ensure that inf-sup stability is maintained. In this paper, we obtain inf-sup stability by using the $Q_2 - P_{-1}$ approximation in which the domain is discretized into shape-regular quadrilaterals on which the velocity space is piecewise continuous biquadratic and the pressure space is piecewise discontinuous linear [34, Chapter 3].

Throughout this paper, we will use the notation \mathbf{M}_H to denote the mass matrix for a finite-dimensional Hilbert space H .

2.3 The stochastic Stokes-Brinkman problem

Let $(\Omega, \Sigma, \mathcal{P})$ be a complete probability space with sample space Ω , σ -algebra Σ , and probability measure \mathcal{P} . We assume stochasticity in the permeability tensor $\mathbf{K}(\omega)$ for $\omega \in \Omega$. For simplicity, we focus only on stochastic permeability, although stochastic viscosity, forcing term, or boundary conditions may be treated in a similar manner.

We assume that ω may be well-approximated by a finite number of random variables $\xi = (\xi_1, \dots, \xi_M)^T$ with $\xi_m \in \Gamma_m \equiv [a_m, b_m]$ and $\xi \in \Gamma \equiv \prod_{m=1}^M \Gamma_m$. This may arise, for example, from a truncated Karhuen-Loeve expansion [11, 12] or from a partitioning of the domain into subdomains. When applying the reduced basis methods discussed in Section 4, it will be convenient if the parametrization admits an affine decomposition for the bilinear form (5) and the linear functional (7) such that

$$\begin{aligned} a(u, v; \xi) &= \sum_{i=1}^{n_A} \theta_i^A(\xi) a_i(u, v) \\ \ell_1(v; \xi) &= \sum_{i=1}^{n_f} \theta_i^f(\xi) f_i(v), \end{aligned}$$

where $\{a_i\}_{i=1}^{n_A}$ are parameter-independent bilinear forms, $\{f_i\}_{i=1}^{n_f}$ are parameter-independent linear functionals, and $\{\theta_i^A\}_{i=1}^{n_A}$ and $\{\theta_i^f\}_{i=1}^{n_f}$ are functions mapping the parameter ξ to the coefficients of the affine decomposition. We note that the bilinear form (6) and the linear functional (8) are independent of the permeability and are, therefore, parameter-independent.

Assuming such an affine decomposition, the saddle-point problem (14) is now parametrized as

$$\begin{bmatrix} \mathbf{A}(\xi) & \mathbf{B}^T \\ \mathbf{B} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{u}(\xi) \\ \mathbf{p}(\xi) \end{bmatrix} = \begin{bmatrix} \mathbf{f}(\xi) \\ \mathbf{g} \end{bmatrix}. \quad (15)$$

The matrix $\mathbf{A}(\xi)$ exhibits an affine decomposition

$$\mathbf{A}(\xi) = \sum_{i=1}^{n_A} \theta_i^A(\xi) \mathbf{A}_i, \quad (16)$$

where $\{\mathbf{A}_i\}_{i=1}^{n_A}$ are discretizations of the parameter-independent bilinear forms $\{a_i\}_{i=1}^{n_A}$. Similarly, $\mathbf{f}(\xi)$ exhibits an affine decomposition

$$\mathbf{f}(\xi) = \sum_{i=1}^{n_f} \theta_i^f(\xi) \mathbf{f}_i, \quad (17)$$

where $\{\mathbf{f}_i\}_{i=1}^{n_f}$ are discretizations of the parameter-independent linear functionals $\{f_i\}_{i=1}^{n_f}$. The affine decompositions (16) and (17) greatly simplify the use of reduced basis methods, as is discussed in Section 4.

3 Adaptive ANOVA and stochastic collocation

In this section, we introduce the ANOVA decomposition [15–17] which is a useful tool for analyzing a multivariate function. In particular, we utilize a truncated ANOVA decomposition to approximate the computationally expensive high-dimensional stochastic problem with a set of cheaper low-dimensional problems. In addition, we adaptively select the most effective ANOVA terms, following the ideas presented in [18, 19].

3.1 ANOVA decomposition

Let $\phi(x; \xi)$ denote a parametrized function for $x \in D$ and $\xi \in \Gamma$. In our application, ϕ may be either the velocity u or pressure p . Let \mathcal{I} denote the index set $\{1, \dots, M\}$. We will refer to a subset $T \subseteq \mathcal{I}$ as a *direction*. ANOVA decomposes the function ϕ into contributions from individual directions T . More formally, the decomposition has the form

$$\phi(x; \xi) = \sum_{T \subseteq \mathcal{I}} \phi_T(x; \xi_T), \quad (18)$$

where ξ_T is a restriction of ξ to the coefficients in T and

$$\phi_T(x; \xi_T) = \int_{\Gamma_{T'}} \phi(x; \xi) d\xi_{T'} - \sum_{S \subset T} \phi_S(x; \xi_S), \quad (19)$$

with $T' \subseteq \mathcal{I}$ denoting the complement of T in \mathcal{I} and $\Gamma_S = \prod_{m \in S} \Gamma_m$ for $S \subseteq \mathcal{I}$. We define the order of a direction T to be its cardinality $|T|$. Thus, the ANOVA decomposition attempts to decompose ϕ into the contributions of each individual direction by removing the contributions from lower-order subdirections. The ANOVA decomposition of ϕ is built by first computing the order-zero term

$$\phi_\emptyset(x) = \int_{\Gamma} \phi(x; \xi) d\xi,$$

then successively computing higher-order terms by first marginalizing out the unused variables with $\int_{\Gamma_{T'}} \phi(x; \xi) d\xi_{T'}$ and then removing the computed effects of previous subdirections with $\sum_{S \subset T} \phi_S(x; \xi_S)$. While the cost of the full ANOVA decomposition is prohibitive in high-dimensional spaces, it may be used to obtain a useful approximation by truncating the expansion by keeping only the low-order terms. Thus, a high-dimensional multivariate function may be approximated through the sum of low-dimensional functions.

The ANOVA decomposition defined by Eqs. 18 and 19 involves the computation of high-dimensional integrals which require expensive Monte Carlo computations. An alternative is anchored ANOVA [18, 19, 25–27] in which an anchor point $c \in \Gamma$ is carefully chosen and the measure in Eq. 19 is replaced by the Dirac measure $\delta(\xi - c)$. The order-zero term may then be computed by evaluating ϕ at the anchor point c :

$$\phi_\emptyset(x) = \phi(x; c).$$

Given the order-zero term, we may compute the first-order terms and then the second-order terms as

$$\begin{aligned} \phi_{\{i\}}(x; \xi_{\{i\}}) &= \phi(x; c, \xi, \{i\}) - \phi_\emptyset(x) \\ \phi_{\{i,j\}}(x; \xi_{\{i,j\}}) &= \phi(x; c, \xi, \{i, j\}) - \phi_{\{i\}}(x; \xi_{\{i\}}) \\ &\quad - \phi_{\{j\}}(x; \xi_{\{j\}}) - \phi_\emptyset(x). \end{aligned}$$

The notation $\phi(x; c, \xi, T)$ is to be understood as evaluating ϕ at the parameter ξ which takes on the values

$$\widehat{\xi}_i = \begin{cases} c_i & i \notin T \\ \xi_i & i \in T. \end{cases}$$

Third and higher-order terms may be computed in a similar fashion.

It was shown in [18] that the mean of ξ often serves as a good choice of anchor point, and we shall use this choice.

3.2 Stochastic collocation

We may utilize the ANOVA decomposition to compute moments of ϕ . The mean of ϕ may simply be obtained by summing the means of the individual terms:

$$\mathbb{E}[\phi(x; \xi)] = \sum_{T \subseteq \mathcal{I}} \mathbb{E}[\phi_T(x; \xi_T)]. \quad (20)$$

A property of the standard ANOVA decomposition is the orthogonality of its terms. Due to this, higher-order moments such as the variance of ϕ may also be obtained by summing the moments of the individual terms. However, the orthogonality property only holds if the same measure is used in both the computation of the ANOVA terms and in the statistical moments. This fails in the case of anchored ANOVA. We remark that, for a good choice of an anchor point, the orthogonality may be approximately preserved. For general choices of anchor points, a method involving the covariances of individual terms is needed for good accuracy [35]. That is, we compute the variance by summing the covariances of all pairs of directions:

$$\begin{aligned} &\mathbb{E}[(\phi(x; \xi) - \mu)^2] \\ &= \sum_{T, S \subseteq \mathcal{I}} \mathbb{E}[(\phi_T(x; \xi_T) - \mu_T)(\phi_S(x; \xi_S) - \mu_S)] \end{aligned}$$

where μ denotes the mean (20) and μ_T denotes the mean of ϕ_T .

To estimate the moments of individual ANOVA terms, we use stochastic collocation. In stochastic collocation, we interpolate the stochastic solution by a set of Lagrange polynomials at collocation points

$$\phi(x; \xi) \approx \sum_{\xi^{(k)} \in \Theta} \phi_c(x; \xi^{(k)}) L_{\xi^{(k)}}(\xi)$$

where $\Theta \subset \Gamma$ is the set of collocation points and $\{L_{\xi^{(k)}}\}$ are the Lagrange polynomials. The coefficients $\phi_c(x; \xi^{(k)})$ are obtained by function evaluations at specific realizations $\xi^{(k)}$ of ξ . Thus, stochastic collocation methods reduce solution of the stochastic problem to a set of deterministic solves at specific sample points.

We select collocation points for each direction T as follows. For each index $i \in T$, a set Θ_i of p_i points and associated weights are selected as nodes according to a quadrature rule on Γ_i . Then the full set of quadrature points is obtained by using a tensor product $\Theta_T = \bigotimes_{i \in T} \Theta_i$. For a quadrature point $\hat{\xi} = \bigotimes_{i \in T} \hat{\xi}_i \in \Theta_T$, the corresponding weight is the product $w(\hat{\xi}) = \prod_{i \in T} w_i(\hat{\xi}_i)$ where $w_i(\hat{\xi}_i)$ denotes the weight of $\hat{\xi}_i$ in the quadrature rule for index i . The collocation means for direction T and for the full function may then be computed as

$$\begin{aligned} \mathbb{E}_{\text{sc}}[\phi_T(x; \xi_T)] &= \sum_{\hat{\xi} \in \Theta_T} \phi_T(x; \hat{\xi}) w(\hat{\xi}) \\ \mathbb{E}_{\text{sc}}[\phi(x; \xi)] &= \sum_{T \subseteq \mathcal{I}} \mathbb{E}[\phi_T(x; \xi_T)]. \end{aligned} \quad (21)$$

For a direction S , let $\tilde{\phi}_S(x; \xi_S) = \phi_S(x; \xi_S) - \mathbb{E}_{\text{sc}}[\phi_S(x; \xi_S)]$. To compute the covariance of S and T , we need quadrature points for $S \cup T$, although we need not form these quadrature points explicitly. We partition $S \cup T$ into three disjoint sets: $S \cap T$, $S \setminus T$, and $T \setminus S$. For $\hat{\xi}_{S \cap T} \in \Theta_{S \cap T}$, define

$$Z_{S \setminus T}(\hat{\xi}_{S \cap T}) = \sum_{\substack{\hat{\xi}_S \in \Theta_S \\ \hat{\xi}_S |_{S \cap T} = \hat{\xi}_{S \cap T}}} w_S(\hat{\xi}_S) \tilde{\phi}_S(x; \hat{\xi}_S).$$

We may then compute the collocation covariances and variances as

$$\begin{aligned} \mathbb{E}_{\text{sc}}[\tilde{\phi}_S(x; \xi_S) \tilde{\phi}_T(x; \xi_T)] &= \sum_{\hat{\xi} \in \Theta_{S \cap T}} \frac{1}{w(\hat{\xi})} Z_{S \setminus T}(\hat{\xi}) Z_{T \setminus S}(\hat{\xi}) \\ \mathbb{E}_{\text{sc}}[(\phi(x; \xi) - \mu)^2] &= \sum_{S, T \subseteq \mathcal{I}} \mathbb{E}_{\text{sc}}[\tilde{\phi}_S(x; \xi_S) \tilde{\phi}_T(x; \xi_T)]. \end{aligned} \quad (22)$$

For truncated ANOVA, we replace (21) and (22) with sums over the used directions.

If a truncated ANOVA decomposition is used to compute (21) and (22), immediate savings over the full tensor product collocation are apparent. Assuming each index

has polynomial order p , the full tensor product collocation requires function evaluations at p^M collocation points, which quickly becomes prohibitive even for moderate M . However, if an ANOVA decomposition truncated at level ℓ is used, then the total number of collocation points is reduced to $\sum_{l=0}^{\ell} \binom{M}{l} p^l$, which is substantially less than p^M .

3.3 Adaptive ANOVA

Instead of simply truncating (18) at a specified level, we could attempt to adaptively select which directions contribute the most significantly. Such an adaptive scheme was developed in [18, 19]. For each direction T , we can score its contribution to the mean using

$$\eta^A(T) = \frac{\|\mathbb{E}_{\text{sc}}[u_T]\|_V + \|\mathbb{E}_{\text{sc}}[p_T]\|_Q}{\|\sum_{|S| < |T|} \mathbb{E}_{\text{sc}}[u_S]\|_V + \|\sum_{|S| < |T|} \mathbb{E}_{\text{sc}}[p_S]\|_Q}. \quad (23)$$

The form of this indicator was used in [29] when applying adaptive ANOVA techniques to the Stokes equations due to the mixed formulation of the model problem. An alternative indicator which uses the variance rather than the mean may be used instead. In this paper, we use the mean.

We consider a direction T *active* if it is included in the ANOVA decomposition. We consider any active direction T whose indicator (23) exceeds a given tolerance ε^A to be *effective*. Let \mathcal{J}_l denote the set of active directions at level l , and let $\tilde{\mathcal{J}}_l \subseteq \mathcal{J}_l$ denote the set of effective directions at level l :

$$\tilde{\mathcal{J}}_l \equiv \{T \in \mathcal{J}_l : \eta^A(T) > \varepsilon^A\}. \quad (24)$$

Then we choose the next-level active directions \mathcal{J}_{l+1} as those level $l+1$ directions T satisfying $S \in \tilde{\mathcal{J}}_l$ for every level l subdirection S of T :

$$\mathcal{J}_{l+1} \equiv \{T : |T| = l+1 \text{ and } S \in \tilde{\mathcal{J}}_l \forall S \subset T, |S| = l\}. \quad (25)$$

Thus, we build level $l+1$ active directions by considering only those directions which can be built from effective level l directions. This heuristic is based upon the idea that, if a direction T is important, then its subdirections at the previous level will likely also be important.

4 Reduced basis methods

Whereas the anchored ANOVA method reduces the number of collocation points needed to accurately estimate the moments, solving (15) at each collocation point may still be prohibitive. In the context of reduced basis methods, we refer to solving (15) as a *high-fidelity* solve. Reduced basis methods [20–23] may be used in conjunction with ANOVA [28, 29] to reduce the number of high-fidelity solves. Suppose that the solutions at the collocation points

may be well-approximated by a low-dimensional space. Significant computational savings may then be obtained by reducing the full computation to this low-dimensional space. We seek spaces $V_r \subset V_h$ and $Q_r \subset Q_h$ of dimensions much smaller than those of V_h and Q_h but for which the solution $(u_r(\xi), p_r(\xi)) \in V_r \times Q_r$ to the variational problem

$$a(u_r(\xi), v_r; \xi) + b(v_r, p_r(\xi)) = \ell_1(v_r; \xi), \quad \forall v_r \in V_r, \quad (26)$$

$$b(u_r(\xi), q_r) = \ell_2(q_r), \quad \forall q_r \in Q_r \quad (27)$$

accurately approximates the solution $(u_h(\xi), p_h(\xi)) \in V_h \times Q_h$ to the high-fidelity problem.

Let \mathbf{V}_r denote a matrix whose columns form a V -orthogonal basis for the space V_r , and, similarly, let \mathbf{Q}_r denote a matrix whose columns form a Q -orthogonal basis for the space Q_r . Then Eqs. 26 and 27 is equivalent to solving the following saddle-point problem as in (15)

$$\begin{bmatrix} \mathbf{V}_r^T \mathbf{A}(\xi) \mathbf{V}_r & \mathbf{V}_r^T \mathbf{B}^T \mathbf{Q}_r \\ \mathbf{Q}_r^T \mathbf{B} \mathbf{V}_r & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{u}_r(\xi) \\ \mathbf{p}_r(\xi) \end{bmatrix} = \begin{bmatrix} \mathbf{V}_r^T \mathbf{f}(\xi) \\ \mathbf{Q}_r^T \mathbf{g} \end{bmatrix}. \quad (28)$$

The approximate solutions in V_h and Q_h may then be obtained as $\mathbf{V}_r \mathbf{u}_r(\xi)$ and $\mathbf{Q}_r \mathbf{p}_r(\xi)$.

Care must be taken to ensure that the reduced basis problem (26) and (27) is inf-sup stable. Whereas the coercivity and continuity of a on the reduced basis space follow directly from the coercivity and continuity of a on the high-fidelity space, the inf-sup condition (11) is not immediately satisfied. The condition may be satisfied by properly enriching the velocity space V_r relative to the pressure space. Let $T: Q \rightarrow V$ denote the *supremizer* operator such that, for a given $q \in Q$, Tq satisfies

$$(Tq, v)_V = b(v, q) \quad \forall v \in V. \quad (29)$$

Then the inf-sup condition is satisfied if, for each $q \in Q_r$, $Tq \in V_r$ [30]. Let $\{u_r^1, \dots, u_r^s\}$ be a set of s high-fidelity velocity solutions from which we seek to build the reduced basis and $\{p_r^1, \dots, p_r^s\}$ be the corresponding set of s high-fidelity pressure solutions. We then obtain an inf-sup stable approximation by defining

$$V_r \equiv \text{span} \{u_r^1, \dots, u_r^s, Tp_r^1, \dots, Tp_r^s\}$$

$$Q_r \equiv \text{span} \{p_r^1, \dots, p_r^s\}.$$

Note that, with this procedure, the dimension of V_r is always twice that of Q_r .

The efficiency of the reduced basis method stems from the dimension of Eq. 28 being significantly smaller than the dimension of Eq. 15. The accuracy relies upon the construction of rigorous a posteriori error estimates which bound the errors of the reduced basis solutions in the V and Q norms. These will be the subject of Section 5. For now, we simply assume such error estimates are available.

We desire the entire computation on the reduced space to be independent of the size of the high-fidelity spaces. For solving (28), this means that the assembly of the saddle-point system must be done independently of the size of the high-fidelity problem. For this, we exploit the affine decomposition (16) and (17) to obtain the reduced basis affine decompositions

$$\begin{aligned} \mathbf{V}_r^T \mathbf{A}(\xi) \mathbf{V}_r &= \sum_{i=1}^{n_A} \boldsymbol{\theta}_i^A(\xi) \mathbf{V}_r^T \mathbf{A}_i \mathbf{V}_r \\ \mathbf{V}_r^T \mathbf{f}(\xi) &= \sum_{i=1}^{n_f} \boldsymbol{\theta}_i^f(\xi) \mathbf{V}_r^T \mathbf{f}_i. \end{aligned}$$

Thus, after constructing the reduced basis, we compute and store the parameter-independent quantities $\mathbf{V}_r^T \mathbf{A}_i \mathbf{V}_r$, $\mathbf{V}_r^T \mathbf{f}_i$, $\mathbf{Q}_r^T \mathbf{B} \mathbf{V}_r$, and $\mathbf{Q}_r^T \mathbf{g}$.

4.1 Constructing the reduced basis

We form the reduced basis using a variation of the standard greedy algorithm [23, Chapter 7] which is summarized in Algorithm 1. In the standard greedy algorithm, the reduced basis is formed in a potentially expensive offline step. A finite training sample set $\Xi \subset \Gamma$ is selected from the space of parameters such that Ξ is a good representation of the parameter space. To fit within our full reduced basis ANOVA algorithm (Algorithm 2), we allow the case where the greedy training algorithm can extend an existing reduced basis. If no reduced basis currently exists, a sample point $\xi \in \Xi$ is selected at which a high-fidelity solve is performed to initialize the reduced basis. For each iteration of training, reduced basis solves are performed over all points in Ξ and the a posteriori error estimates are computed for each reduced basis solution. If all error estimates are below a prescribed tolerance ε^{RB} , then training stops and the current reduced basis is output. Otherwise, the parameter which attains the largest error estimate is selected and a high-fidelity solve is performed at this parameter. The high-fidelity solution is then added to the reduced basis. This process repeats until all error estimates are below the desired tolerance. To ensure inf-sup stability, in addition to adding the high-fidelity velocity solution to the velocity reduced basis, we must also add the application of the supremizer (29) to the corresponding pressure solution. A Gram-Schmidt orthonormalization procedure should be used to ensure that the reduced basis matrices remain orthonormal with respect to the appropriate inner product.

Our variation of the standard greedy algorithm follows that of [36] for reduced basis collocation in which the formation of the collocation points and the generation of the reduced basis are not separated. That is, as we are building the collocation points at each ANOVA level, we

Algorithm 1 Greedy algorithm for generating a reduced basis using a training set \mathcal{E} and a tolerance ε^{RB} .

```

1: if  $\mathbf{V}_r$  and  $\mathbf{Q}_r$  are uninitialized then
2:   choose  $\hat{\xi} \in \mathcal{E}$ 
3:   compute high-fidelity solution (15)  $(\mathbf{u}(\hat{\xi}), \mathbf{p}(\hat{\xi}))$ 
4:   compute supremizer (29)  $\mathbf{v}(\hat{\xi}) = \mathbf{M}_V^{-1} \mathbf{B}^T \mathbf{p}(\hat{\xi})$ 
5:    $\mathbf{V}_r \leftarrow [\mathbf{u}(\hat{\xi}) \mid \mathbf{v}(\hat{\xi})]$ 
6:    $\mathbf{Q}_r \leftarrow [\mathbf{p}(\hat{\xi})]$ 
7: end if
8: loop
9:   for  $\xi \in \mathcal{E}$  do
10:    compute RB solution (28)  $(\mathbf{u}_r(\xi), \mathbf{p}_r(\xi))$ 
11:    compute a posteriori error estimate (34)  $\Delta_r(\xi)$ 
12:   end for
13:   let  $\Delta = \max_{\xi \in \mathcal{E}} \Delta_r(\xi)$ 
14:   if  $\Delta < \varepsilon^{\text{RB}}$  then
15:     return reduced basis and offline
        quantities
16:   else
17:     let  $\hat{\xi} = \operatorname{argmax}_{\xi \in \mathcal{E}} \Delta_r(\xi)$ 
18:     compute high-fidelity
        solutions (15)  $(\mathbf{u}(\hat{\xi}), \mathbf{p}(\hat{\xi}))$ 
19:     compute supremizer (29)  $\mathbf{v}(\hat{\xi}) = \mathbf{M}_V^{-1} \mathbf{B}^T \mathbf{p}(\hat{\xi})$ 
20:      $\mathbf{V}_r \leftarrow [\mathbf{V}_r \mid \mathbf{u}(\hat{\xi}) \mid \mathbf{v}(\hat{\xi})]$ 
21:      $\mathbf{Q}_r \leftarrow [\mathbf{Q}_r \mid \mathbf{p}(\hat{\xi})]$ 
22:   end if
23: end loop

```

use those collocation points as our training set \mathcal{E} . In [36], the authors proposed a single sweep through these points by augmenting the reduced basis whenever a single parameter's error estimate exceeds the tolerance. In this paper, however, we use the standard approach of computing error estimates over all training points and augmenting the reduced basis with a high-fidelity solve at the point with the largest error estimate. We remark that this approach is easily parallelized as the reduced basis solves and error estimate computations are embarrassingly parallel. The approach in [36], however, is better suited for serial computation. The full reduced basis ANOVA algorithm is summarized in Algorithm 2.

5 A posteriori error estimates

Essential for the successful application of reduced basis methods is the computation of rigorous a posteriori error estimates. We follow the approach of [31] for saddle-point problems, which develops residual-based error estimates based upon the Brezzi stability theory. Let $u_r(\xi) \in V_r$ and $p_r(\xi) \in Q_r$ denote the reduced basis velocity and pressure solutions, respectively, for parameter ξ , and let $u_h(\xi) \in V_h$ and $p_h(\xi) \in Q_h$ denote the high-fidelity

Algorithm 2 Reduced Basis Adaptive ANOVA with initial level ℓ , maximum level L , reduced basis tolerance ε^{RB} , and ANOVA tolerance ε^{A} .

```

1: compute anchor point  $c$  as mean of distribution
2: solve  $\mathbf{u}(c), \mathbf{p}(c)$  of high-fidelity problem (15) at anchor
   point  $c$ 
3: compute supremizer (29)  $\mathbf{v}(c) = \mathbf{M}_V^{-1} \mathbf{B}^T \mathbf{p}(c)$ 
4:  $\mathbf{V}_r \leftarrow [\mathbf{u}(c) \mid \mathbf{v}(c)]$ 
5:  $\mathbf{Q}_r \leftarrow [\mathbf{p}(c)]$ 
6: initialize set of active directions  $\mathcal{J}_\ell$  as all directions of
   order up to  $\ell$ 
7: initialize  $\mathcal{E}_0 \leftarrow \{c\}$ 
8: for  $l = \ell$  to  $L$  do
9:   generate all collocation points  $\mathcal{E}_l$  for the active
   directions  $\mathcal{J}_l$ 
10:  form reduced basis training set  $\mathcal{E} \leftarrow \mathcal{E}_l \setminus \mathcal{E}_0$ 
11:  update reduced basis matrices  $\mathbf{V}_r$  and  $\mathbf{Q}_r$  using
   training set  $\mathcal{E}$  and tolerance  $\varepsilon^{\text{RB}}$  (see Algorithm 1)
12:   $\mathcal{E}_0 \leftarrow \mathcal{E}_l \cup \mathcal{E}_0$ 
13:  perform reduced basis solves on  $\mathcal{E}_0$  and compute
   ANOVA indicators (23)
14:  if  $l < L$  then
15:    compute effective directions  $\widetilde{\mathcal{J}}_l$  (24)
16:    compute next-level active directions  $\mathcal{J}_{l+1}$  (25)
17:    if  $|\mathcal{J}_{l+1}| = 0$  then
18:      return reduced basis and ANOVA solutions
19:    end if
20:  end if
21: end for
22: return reduced basis and ANOVA solutions

```

velocity and pressure solutions, respectively. We are interested in bounding the velocity and pressure errors

$$e_r^u(\xi) \equiv u_h(\xi) - u_r(\xi) \in V_h$$

$$e_r^p(\xi) \equiv p_h(\xi) - p_r(\xi) \in Q_h.$$

We note that these errors are the solution to the Stokes-Brinkman problem

$$a(e_r^u(\xi), v; \xi) + b(v, e_r^p(\xi)) = r_1(v; \xi), \quad \forall v \in V_h \quad (30)$$

$$b(e_r^u(\xi), q) = r_2(q; \xi), \quad \forall q \in Q_h \quad (31)$$

where $r_1(\cdot; \xi) \in V_h^*$ and $r_2(\cdot; \xi) \in Q_h^*$ are the reduced basis residuals defined by

$$r_1(v; \xi) \equiv \ell_1(v; \xi) - a(u_r(\xi), v; \xi) - b(v, p_r(\xi)) \quad (32)$$

$$r_2(q; \xi) \equiv \ell_2(q) - b(u_r(\xi), q). \quad (33)$$

We are interested in tight upper bounds $\Delta_r^u(\xi)$ and $\Delta_r^p(\xi)$ which satisfy

$$\|e_r^u(\xi)\|_V \leq \Delta_r^u(\xi)$$

$$\|e_r^p(\xi)\|_Q \leq \Delta_r^p(\xi).$$

We may then form an a posteriori error estimate over the whole solution as

$$\Delta_r(\xi) = \sqrt{\Delta_r^u(\xi)^2 + \Delta_r^p(\xi)^2}. \quad (34)$$

Applying the stability bounds (12) and (13) from Theorem 1 to Eqs. 30 and 31, we obtain the following upper bounds for the errors:

$$\|e_r^u(\xi)\|_V \leq \frac{1}{\alpha(\xi)} \|r_1(\cdot; \xi)\|_{V^*} + \frac{2}{\beta} \sqrt{\frac{\gamma(\xi)}{\alpha(\xi)}} \|r_2(\cdot; \xi)\|_{Q^*} \quad (35)$$

$$\|e_r^p(\xi)\|_Q \leq \frac{2}{\beta} \sqrt{\frac{\gamma(\xi)}{\alpha(\xi)}} \|r_1(\cdot; \xi)\|_{V^*} + \frac{\gamma(\xi)}{\beta^2} \|r_2(\cdot; \xi)\|_{Q^*}, \quad (36)$$

where $\alpha(\xi)$ is the coercivity constant

$$\alpha(\xi) \equiv \inf_{v \in V_h} \frac{a(v, v; \xi)}{\|v\|_V^2},$$

$\gamma(\xi)$ is the continuity constant

$$\gamma(\xi) \equiv \sup_{u, v \in V_h} \frac{a(u, v; \xi)}{\|u\|_V \|v\|_V},$$

and β is the inf-sup constant

$$\beta \equiv \inf_{q \in Q_h} \sup_{v \in V_h} \frac{b(v, q)}{\|q\|_Q \|v\|_V}.$$

Note that the inf-sup constant β is defined independently of the parameter ξ . Thus, we only need to solve a single eigenvalue problem to determine its value over all parameters. The coercivity $\alpha(\xi)$ and continuity $\gamma(\xi)$ constants are parameter-dependent so that computing their exact values would require solving separate eigenvalue problems for each parameter. To reduce the computational cost, we instead compute tight bounds in an efficient manner that is independent of the dimension of V_h . Details of this procedure are described in Section 5.2.

Using a lower bound $\alpha^{\text{LB}}(\xi) \leq \alpha(\xi)$ for the coercivity constant and an upper bound $\gamma^{\text{UB}}(\xi) \geq \gamma(\xi)$, for the continuity constant, we may define our error estimates as

$$\Delta_r^u(\xi) \equiv \frac{1}{\alpha^{\text{LB}}(\xi)} \|r_1(\cdot; \xi)\|_{V^*} + \frac{2}{\beta} \sqrt{\frac{\gamma^{\text{UB}}(\xi)}{\alpha^{\text{LB}}(\xi)}} \|r_2(\cdot; \xi)\|_{Q^*} \quad (37)$$

$$\Delta_r^p(\xi) \equiv \frac{2}{\beta} \sqrt{\frac{\gamma^{\text{UB}}(\xi)}{\alpha^{\text{LB}}(\xi)}} \|r_1(\cdot; \xi)\|_{V^*} + \frac{\gamma^{\text{UB}}(\xi)}{\beta^2} \|r_2(\cdot; \xi)\|_{Q^*}. \quad (38)$$

Since these are upper bounds for the stability estimates (35) and (36), we observe that these are true upper bounds for the reduced basis errors.

5.1 Efficient computation of the dual norm

Computing the error estimates (37) and (38) requires computation of the dual norms of the residuals. Dual norm computations are best performed through use of the Riesz representation. Let $R_H: H \rightarrow H^*$ denote the Riesz isomorphism from a Hilbert space H to its dual defined

by $\langle R_H u, v \rangle_H = (u, v)_H$ for all $u, v \in H$. Here, $\langle \cdot, \cdot \rangle_H$ denotes the duality pairing and $(\cdot, \cdot)_H$ denotes the inner product. Then, by the Riesz representation theorem, for any $f \in H^*$, we have $\|f\|_{H^*} = \|R_H^{-1} f\|_H$. For finite-dimensional spaces, it is apparent that the operator R_H is represented by the mass matrix for the H -norm. Thus, if \mathbf{f} denotes the vector of coefficients for $f \in H^*$, then $\|f\|_{H^*}^2 = \mathbf{f}^T \mathbf{M}_H^{-1} \mathbf{f}$.

In order for the error estimates to be computed efficiently during the online phase, we desire the cost to be independent of the dimension of the finite element spaces V_h and Q_h . To this effect, we utilize the affine decompositions (16) and (17). The dual norm of the residual (32) is then computed as

$$\begin{aligned} \|r_1(\cdot; \xi)\|_{V^*}^2 &= \sum_{i=1}^{n_f} \sum_{j=1}^{n_f} \theta_i^f(\xi) \theta_j^f(\xi) \mathbf{C}_{ij} \\ &\quad - 2 \left(\sum_{i=1}^{n_f} \sum_{j=1}^{n_A} \theta_i^f(\xi) \theta_j^A(\xi) \mathbf{D}_{ij} \right) \mathbf{u}_r \\ &\quad - 2 \left(\sum_{i=1}^{n_f} \theta_i^f(\xi) \mathbf{E}_i \right) \mathbf{p}_r \\ &\quad + \mathbf{u}_r^T \left(\sum_{i=1}^{n_A} \sum_{j=1}^{n_A} \theta_i^A(\xi) \theta_j^A(\xi) \mathbf{F}_{ij} \right) \mathbf{u}_r \\ &\quad + 2 \mathbf{u}_r^T \left(\sum_{i=1}^{n_A} \theta_i^A(\xi) \mathbf{G}_i \right) \mathbf{p}_r \\ &\quad + \mathbf{p}_r^T \mathbf{H} \mathbf{p}_r, \end{aligned} \quad (39)$$

where

$$\mathbf{C}_{ij} = \mathbf{f}_i^T \mathbf{M}_V^{-1} \mathbf{f}_j$$

$$\mathbf{D}_{ij} = \mathbf{f}_i^T \mathbf{M}_V^{-1} \mathbf{A}_j \mathbf{V}_r$$

$$\mathbf{E}_i = \mathbf{f}_i^T \mathbf{M}_V^{-1} \mathbf{B}^T \mathbf{Q}_r$$

$$\mathbf{F}_{ij} = \mathbf{V}_r^T \mathbf{A}_i^T \mathbf{M}_V^{-1} \mathbf{A}_j \mathbf{V}_r$$

$$\mathbf{G}_i = \mathbf{V}_r^T \mathbf{A}_i \mathbf{M}_V^{-1} \mathbf{B}^T \mathbf{Q}_r$$

$$\mathbf{H} = \mathbf{Q}_r^T \mathbf{B} \mathbf{M}_V^{-1} \mathbf{B}^T \mathbf{Q}_r,$$

and the dual norm of the residual (33) is computed as

$$\|r_2(\cdot; \xi)\|_{Q^*}^2 = \mathbf{R} - 2 \mathbf{S} \mathbf{u}_r + \mathbf{u}_r^T \mathbf{T} \mathbf{u}_r$$

where

$$\mathbf{R} = \mathbf{g}^T \mathbf{M}_Q^{-1} \mathbf{g}$$

$$\mathbf{S} = \mathbf{g}^T \mathbf{M}_Q^{-1} \mathbf{B} \mathbf{V}_r$$

$$\mathbf{T} = \mathbf{V}_r^T \mathbf{B}^T \mathbf{M}_Q^{-1} \mathbf{B} \mathbf{V}_r.$$

The terms \mathbf{C}_{ij} , \mathbf{D}_{ij} , \mathbf{E}_i , \mathbf{F}_{ij} , \mathbf{G}_i , \mathbf{H} , \mathbf{R} , \mathbf{S} , \mathbf{T} are all parameter-independent and can be computed and stored during the offline phase. Thus, the dual norms for the residuals may be computed efficiently during the online phase in a cost that is independent of the dimensions of V_h and Q_h .

Remark 2 For moderately large reduced bases, the offline quantities may consume a significant amount of memory. In particular, there are n_A^2 matrices of the form \mathbf{F}_{ij} . If the size of the velocity reduced basis is N , then each of these matrices will contain N^2 entries, for a storage requirement of $n_A^2 N^2$ double-precision floating-point numbers. Nonetheless, the computation is still independent of the size of the high-fidelity problem. This, however, should be considered when evaluating the efficiency of the reduced basis method.

5.2 Efficient computation of stability bounds using SCM

As previously remarked, we seek efficiently computable bounds on the coercivity and continuity constants with cost that is independent of the high-fidelity problem. A popular method for obtaining such bounds is the Successive Constraint Method (SCM), first proposed in [37] and subsequently refined in [38]. We recall here the method as presented in [38].

For simplicity, we focus on the coercivity constant, since the method may easily be extended to the continuity case. Recall that the coercivity constant is the largest $\alpha(\xi) > 0$ such that $a(u, u; \xi) \geq \alpha(\xi) \|u\|_V$ for all $u \in V$. Using the affine decomposition of a , we may write

$$\alpha(\xi) = \inf_{u \in V} \sum_{i=1}^{n_A} \theta_i^A(\xi) \frac{a_i(u, u)}{\|u\|_V}. \quad (40)$$

Defining the set

$$\begin{aligned} \mathcal{Y} &\equiv \{y = (y_1, \dots, y_{n_A})^T \in \mathbb{R}^{n_A} : \\ &\quad \exists u \in V \ni y_i = \frac{a_i(u, u)}{\|u\|_V}, i = 1, \dots, n_A\}, \end{aligned}$$

we may express (40) as the minimization problem

$$\alpha(\xi) = \inf_{y \in \mathcal{Y}} \sum_{i=1}^{n_A} \theta_i^A(\xi) y_i.$$

We may obtain a lower bound for $\alpha(\xi)$ by replacing \mathcal{Y} with a set $\mathcal{Y}_{LB}(\xi)$ such that $\mathcal{Y} \subset \mathcal{Y}_{LB}(\xi)$. By choosing $\mathcal{Y}_{LB}(\xi)$ to be a convex polyhedron, we obtain the following linear programming problem

$$\alpha^{LB}(\xi) = \min_{y \in \mathcal{Y}_{LB}(\xi)} \sum_{i=1}^{n_A} \theta_i^A(\xi) y_i.$$

To define $\mathcal{Y}_{LB}(\xi)$, we first note that the box \mathcal{B} defined as

$$\begin{aligned} \mathcal{B} &\equiv \{y \in \mathbb{R}^{n_A} : \\ &\quad \inf_{u \in V} \frac{a_i(u, u)}{\|u\|_V} \leq y_i \leq \sup_{u \in V} \frac{a_i(u, u)}{\|u\|_V}, i = 1, \dots, n_A\} \end{aligned}$$

clearly satisfies $\mathcal{Y} \subset \mathcal{B}$. Forming \mathcal{B} involves the solution of $2n_A$ eigenvalue problems. Suppose we are given two

sets of parameters— \mathcal{E}_E for which exact coercivity constants have been computed and \mathcal{E}_P for which lower bounds have been computed. Then a finer $\mathcal{Y}_{LB}(\xi)$ may be obtained as follows. First, choose $M_E, M_P \geq 0$ and locate the M_E points $\{\xi_E^1, \dots, \xi_E^{M_E}\}$ in \mathcal{E}_E that are closest to ξ and the M_P points $\{\xi_P^1, \dots, \xi_P^{M_P}\}$ in \mathcal{E}_P that are closest to ξ . We then include the following $M_E + M_P$ inequality constraints in addition to those defined by the box \mathcal{B} :

$$\begin{aligned} \sum_{i=1}^{n_A} \theta_i^A(\xi_E^j) y_i &\geq \alpha(\xi_E^j), j = 1, \dots, M_E \\ \sum_{i=1}^{n_A} \theta_i^A(\xi_P^j) y_i &\geq \alpha^{LB}(\xi_P^j), j = 1, \dots, M_P. \end{aligned}$$

It is clear that the set $\mathcal{Y}_{LB}(\xi)$ is a superset of \mathcal{Y} . We thus reduce the problem of computing $\alpha^{LB}(\xi)$ to solving a linear program in n_A variables with $2n_A + M_E + M_P$ linear inequality constraints, which is independent of the dimension of V_h , as desired.

What remains is to demonstrate how \mathcal{E}_E is selected and the lower bounds for the points in \mathcal{E}_P are obtained. These may be built through an offline training procedure. This procedure will make use of the following upper bound for $\alpha(\xi)$. Given any ξ' for which $\alpha(\xi')$ has been computed, we may use the corresponding eigenvector u' and compute $y' = (y_1, \dots, y_{n_A})^T \in \mathcal{Y}$ where $y_i = a_i(u', u')/\|u'\|_V$. Collecting all such y' into a set \mathcal{Y}_{UB} , we obtain an upper bound as

$$\alpha^{UB}(\xi) = \min_{y \in \mathcal{Y}_{UB}} \sum_{i=1}^{n_A} \theta_i^A(\xi) y_i.$$

Since \mathcal{Y}_{UB} is small and finite, we may compute this upper bound efficiently by simply enumerating over all elements of \mathcal{Y}_{UB} . Furthermore, since clearly $\mathcal{Y}_{UB} \subset \mathcal{Y}$, we have a true upper bound. Given upper and lower bounds for the coercivity constant, we define the indicator

$$\eta^{SCM}(\xi) \equiv 1 - \frac{\alpha^{LB}(\xi)}{\alpha^{UB}(\xi)} \quad (41)$$

which measures the relative gap between the upper and lower bounds.

Having determined an indicator for the quality of the SCM approximation, we may now describe the offline training algorithm. We begin with a training set \mathcal{E} which should be representative of the parameter set. Choose $\xi' \in \mathcal{E}$, initialize $\mathcal{E}_E = \{\xi'\}$ and $\mathcal{E}_P = \mathcal{E} \setminus \{\xi'\}$, and set $\alpha^{LB}(\xi) = 0$ for all $\xi \in \mathcal{E}$. Then loop as follows. Solve for the exact coercivity constant for ξ' and compute lower bounds using SCM and indicators for all $\xi \in \mathcal{E}_P$. If the largest indicator is below a desired tolerance ϵ^{SCM} , terminate. Otherwise, select a new point $\xi' \in \mathcal{E}_P$ with the largest indicator, compute and store $\alpha(\xi')$, move ξ' from \mathcal{E}_P to \mathcal{E}_E , and update the lower bounds for points

in \mathcal{E}_P using SCM. Repeat until termination. This procedure is summarized in Algorithm 3.

Algorithm 3 Greedy training algorithm for SCM given training set \mathcal{E} .

```

1:  $\mathcal{E}_E \leftarrow \emptyset, \mathcal{E}_P \leftarrow \mathcal{E}$ 
2: choose  $\xi' \in \mathcal{E}$ 
3:  $\alpha^{\text{LB}}(\xi) \leftarrow 0$  for all  $\xi \in \mathcal{E}$ 
4: loop
5:   compute and store  $\alpha(\xi')$ 
6:    $\mathcal{E}_E \leftarrow \mathcal{E}_E \cup \{\xi'\}, \mathcal{E}_P \leftarrow \mathcal{E}_P \setminus \{\xi'\}$ 
7:   if  $\mathcal{E}_P = \emptyset$  then
8:     return SCM data
9:   end if
10:  for  $\xi \in \mathcal{E}_P$  do
11:    update  $\alpha^{\text{LB}}(\xi)$  using SCM
12:    compute indicator  $\eta^{\text{SCM}}(\xi)$  (41)
13:  end for
14:   $\eta \leftarrow \max_{\xi \in \mathcal{E}_P} \eta^{\text{SCM}}(\xi)$ 
15:  if  $\eta < \epsilon^{\text{SCM}}$  then
16:    return SCM data
17:  end if
18:   $\xi' \leftarrow \operatorname{argmax}_{\xi \in \mathcal{E}_P} \eta^{\text{SCM}}(\xi)$ 
19: end loop

```

6 Numerical experiments

We compare the effectiveness of reduced basis ANOVA on a set of three model problems. These model problems are introduced in Section 6.1 and consider both isotropic and anisotropic permeabilities. In Section 6.2, we analyze the performance of the SCM method on the three problems and provide details on the necessary eigenvalue computations. Finally, in Section 6.3, we study the performance of reduced basis ANOVA.

6.1 Model problems

We consider three model problems. All three make use of the simple 2D square domain pictured in Fig. 1. The top and bottom boundaries are designated as walls through which no flow occurs. The left boundary is a parabolic inflow boundary. The right boundary is a do-nothing out-flow boundary. All three problems partition this domain into $n \times n$ uniform subdomains in which permeabilities are constant. All three problems use constant viscosity $\nu = \nu^* = 10^{-3}$ and a zero forcing term.

For the first problem, labeled *iso* in the following discussion, we consider isotropic flow in which the permeability tensor in each subdomain is of the form $k\mathbf{I}$ for a scalar $k > 0$. In this case, we partition the domain into 9×9




Fig. 1 Domain in 2D used for numerical experiments

subdomains. Each subdomain is then partitioned into 12×12 regular quadrilateral elements to form a 108×108 grid. The permeability coefficient of each subdomain is uniformly distributed in an interval $[a, b]$ chosen as follows. First, for each subdomain, a value c is sampled from a $\text{Beta}(0.5, 0.5)$ distribution and then mapped to the interval $[-6, -3]$. A value r is uniformly sampled in the interval $[0.05, 0.15]$, and the chosen interval is $[(1 - r) \cdot 10^c, (1 + r) \cdot 10^c]$. The parameters of the beta distribution were chosen to skew the permeabilities to slightly favor a mix of high permeability regions (near 10^{-3}) and low permeability regions (near 10^{-6}). The mean permeabilities for each subdomain are displayed in Fig. 2 on a log scale.

The resulting stochastic space is then of dimension $N = 81$. The parametrization coincides with the permeabilities in each subdomain, i.e., $\xi \in \mathbb{R}^{81}$ with ξ_i equal to the permeability in subdomain i . The affine decomposition (16) consists of $n_A = N + 1 = 82$ parameter independent matrices as follows:

$$A_i = \begin{cases} A_S & i = 1 \\ A_D^{i-1} & i > 1, \end{cases} \quad (42)$$




Fig. 2 Mean permeabilities for the *iso* problem on a log scale

where A_S is the discretization of the Stokes bilinear form (3), and A_D^k is the discretization of the Darcy bilinear form (4) with support only on subdomain k . The mappings $\{\theta_i^A(\xi)\}_{i=1}^{n_A}$ then take the form

$$\theta_i^A(\xi) = \begin{cases} 1 & i = 1 \\ 1/\xi_{i-1} & i > 1. \end{cases} \quad (43)$$

To apply the Dirichlet boundary conditions, we use the finite element function which interpolates the Dirichlet conditions on the Dirichlet boundary and has zero value on all remaining nodes. Let w_h denote this finite element function. Then, since we define zero forcing term and do nothing Neumann conditions, the discretizations of Eqs. 7 and 8 are simply

$$\ell_1(v_h; \xi) = -a(w_h, v_h; \xi) \quad \forall v_h \in V_h \quad (44)$$

$$\ell_2(q_h) = -b(w_h, q_h) \quad \forall q_h \in Q_h. \quad (45)$$

Note that, since b is parameter-independent and the Dirichlet boundary conditions are parameter-independent, the discretization (45) is parameter-independent. However, due to the dependence of a on the parameter, the discretization (44) is parameter-dependent. It admits an affine decomposition of the form (42) and (43). We note, however, that our choice of w_h suggests that we need only $n_f = 10$ parameter-independent vectors, corresponding to the vector arising from the Stokes bilinear form and the 9 Darcy bilinear forms on the subdomains which border the in-flow boundary.

For the second and third problems, we consider anisotropic flow in which the permeability tensor in each subdomain is of the form $\text{diag}(k_x, k_y)$ for positive scalars $k_x \neq k_y$. The second problem, labeled *aniso1*, considers the case where $k_x < k_y$, that is, where vertical flow is favored. The third problem, labeled *aniso2*, considers the case where $k_x > k_y$, that is, where horizontal flow is favored. Both problems are partitioned into 6 × 6 subdomains, each subdomain partitioned into 18 × 18 elements. This partitioning is to make the size of the problems similar to that of the isotropic problem. Here, the total number of elements is 108 × 108, matching the isotropic problem exactly; and the number of parameters is 72, as there are two parameters per subdomain.

The distributions on permeabilities are chosen in a manner similar to the isotropic problem. The main difference is that the smaller permeability (k_x in the case of *aniso1*, k_y in the case of *aniso2*) has the beta random variable mapped to $[-6, -4.75]$, and the larger permeability (k_y in the case of *aniso1*, k_x in the case of *aniso2*) has the beta random variable mapped to $[-4.25, -3]$. The mean permeabilities for each subdomain and direction are displayed in Figs. 3 and 4 on a log scale.

The anisotropic problems admit a similar affine decomposition to the isotropic problem. However, now there are

two Darcy parameter-independent matrices per subdomain, each discretizations of Eq. 4 with support only from basis functions on the given subdomain and direction. Thus, $n_A = 73$. Furthermore, we have $n_f = 7$, as our choice of w_h only admits non-zero parameter-independent vectors arising from the Stokes bilinear form and then one for each of the 6 subdomains bordering the in-flow boundary (as the Dirichlet in-flow conditions are zero in the y -direction).

Table 1 compares the sizes of the finite element discretizations of each problem. As can be observed, all three problems have the same dimension for the high-fidelity space, suggesting they should all be of similar difficulty. Condition number estimates, computed using the `condest` function in MATLAB, for the assembled matrices when choosing the mean of each distribution as the parameter are presented in Table 2. Here, we observe that all three problems are fairly ill-conditioned. We note that the sizes of the high-fidelity problems were chosen to be small enough that we could perform the high-fidelity solves efficiently using a direct solver. The log permeabilities were also chosen to be within the interval $[-6, -3]$ so as to keep the condition numbers moderate. These condition numbers have an effect on the sharpness of the error bounds, as will be discussed in Section 6.3. In short, larger condition numbers imply that the error bounds may not be as sharp, leading to potentially larger reduced bases than necessary.

Table 3 compares the sizes of the stochastic dimension and number of parameter-independent components in the affine decompositions. The isotropic problem is larger in this case due to the use of more subdomains.

6.2 SCM

For the eigenvalue problems solved during the SCM training (see Algorithm 3), we used the LOBPCG method [39] as implemented in the BLOPEX MATLAB package [40]. The inf-sup constant, being parameter-independent, required only a one-time eigenvalue computation. The specific eigenvalue problem solved was to find the smallest eigenvalue λ of

$$\mathbf{B}\mathbf{M}_V^{-1}\mathbf{B}^T x = \lambda \mathbf{M}_Q x \quad (46)$$

and choosing $\beta = \sqrt{\lambda}$. No preconditioner was used. The coercivity and continuity constants are parameter-dependent and require solving for the smallest and largest, respectively, eigenvalues $\lambda(\xi)$ of the generalized eigenvalue problem

$$\mathbf{A}(\xi)x = \lambda(\xi)\mathbf{M}_V x. \quad (47)$$

For the coercivity constant, we solved for the smallest eigenvalue $\lambda(\xi)$ of Eq. 47 using an incomplete Cholesky factorization of $\mathbf{A}(\xi)$ as a preconditioner and set $\alpha(\xi) = \lambda(\xi)$. Computing the continuity constant requires solving for the largest eigenvalue of Eq. 47. For this, we transformed




Fig. 3 Mean permeabilities for the *aniso1* problem on a log scale




Fig. 4 Mean permeabilities for the *aniso2* problem on a log scale

Table 1 Comparison of sizes of the finite element discretization for each model problem

Problem	Subdomains	Elements	Total elements	Velocity dof	Pressure dof	Total dof
<i>iso</i>	9×9	12×12	11664	92880	34992	127872
<i>aniso1</i>	6×6	18×18	11664	92880	34992	127872
<i>aniso2</i>	6×6	18×18	11664	92880	34992	127872

Table 2 Comparison of condition number estimates for the assembled matrices with the mean parameter

Problem	Condition number
<i>iso</i>	2.6894×10^4
<i>aniso1</i>	5.0727×10^4
<i>aniso2</i>	4.0431×10^3

the problem into solving for the smallest eigenvalue $\mu(\xi)$ of the generalized eigenvalue problem

$$\mathbf{M}_V = \mu(\xi)\mathbf{A}(\xi)x. \quad (48)$$

The continuity constant is then chosen as $\gamma(\xi) = 1/\mu(\xi)$. All LOBPCG computations were performed with a maximum of 1000 iterations and a tolerance of 10^{-10} .

We remark that inverting \mathbf{M}_V is performed during the computation of Eq. 39, computation of Eq. 46, and application of the supremizer (29). For small enough problems, we may compute and store the sparse Cholesky factorization of \mathbf{M}_V . In these cases, we may use these factors to perform an exact solve as a preconditioner for Eq. 48. If a preconditioned iterative method is required for inverting \mathbf{M}_V , then we may use the same preconditioner for Eq. 48. In this paper, we used the Cholesky factors for the preconditioner.

Each application of SCM requires solving a linear programming problem. For this, we used the `linprog` function in MATLAB with the dual-simplex algorithm.

For the SCM training of each model problem, we generated a Halton set of size 50000. We chose $M_E = M_P = 100$ and the tolerance $\epsilon^{SCM} = 0.1$. Table 4 summarizes the number of iterations required for each model problem to attain the prescribed tolerance. As can be observed, the training terminates fairly quickly in all three cases, with the coercivity constant being slightly harder than continuity and the isotropic case being slightly harder than the anisotropic cases. The isotropic case being more difficult is likely due to the larger number of parameters. Figures 5 and 6 display the largest indicator (41) over each iteration of training. The indicators for the *iso* problem are plotted with the solid blue line with circles; the indicators for the *aniso1* problem are plotted with the dashed red line with squares; the indicators for the *aniso2* problem are plotted with the dotted black line with triangles.

Table 3 Comparison of parametrization sizes for each model problem


Problem	Stochastic dimension	n_A	n_f
<i>iso</i>	81	82	10
<i>aniso1</i>	72	73	7
<i>aniso2</i>	72	73	7

Table 4 Number of SCM training iterations to obtain a tolerance of $\epsilon^{SCM} = 0.1$ with a training size of 50000 points and parameters $M_E = M_P = 100$

Problem	Coercivity	Continuity
<i>iso</i>	23	7
<i>aniso1</i>	11	3
<i>aniso2</i>	12	3

6.3 Reduced basis ANOVA

To study the effectiveness of reduced basis ANOVA on the three model problems, we compute reduced basis ANOVA approximations using the following tensor product of reduced basis and ANOVA tolerances. For the reduced basis tolerances, we consider the set $\{1, 0.1, 0.01\}$. For the ANOVA tolerances, we consider the set $\{10^{-4}, 10^{-5}, 10^{-6}\}$. Thus, we performed a total of 9 reduced basis ANOVA computations for each problem. In all problems, we used all ANOVA directions up to level 1 and adaptively selected directions for higher levels. In all problems, the adaptive selection of ANOVA directions terminated at level 2. We used Gauss-Legendre quadrature with polynomial order 5 for generating collocation points. We compare each reduced basis ANOVA computation with the results from a Monte Carlo simulation performed using 10^6 points generated from a Halton set. We assume that the moments computed using the Monte Carlo are accurate enough to be used as the true moments so that we may use it to approximate the error in the moments from the reduced basis ANOVA computations.

**Fig. 5** Largest indicators for coercivity during each iteration of SCM training using a tolerance of $\epsilon^{SCM} = 0.1$ with a training size of 50000 points and parameters $M_E = M_P = 100$




Fig. 6 Largest indicators for continuity during each iteration of SCM training using a tolerance of $\varepsilon^{\text{SCM}} = 0.1$ with a training size of 50000 points and parameters $M_E = M_P = 100$

Table 5 summarizes the L^2 norms for the errors in mean and variance for the *iso* problem, where the error here is taken as the difference between the reduced basis ANOVA approximation and the Monte Carlo simulation. These values are normalized by the norm of the Monte Carlo approximation so as to provide relative errors. The table shows the errors in the velocity, pressure, and combined mean and variances. It is clear that, as the reduced basis tolerance is reduced, better approximations to the mean and variances are obtained. A reduced basis tolerance of 10^{-2} results in excellent approximations even for an ANOVA tolerance of 10^{-4} . However, there appear to be little gains from reducing the ANOVA tolerance. Accuracy improves by increasing the tolerance from 10^{-4} to 10^{-5} ; however, increasing the tolerance from 10^{-5} to 10^{-6} yields no improvement. Tables 6 and 7 present the same information for the *aniso1* and *aniso2* problems, respectively. In these cases, we observe the same trend of increased accuracy when reducing the reduced basis tolerance and stagnating accuracy when reducing the ANOVA tolerance.

The errors over the physical domain D are displayed in Figs. 7 (pressure), 8 (velocity magnitude), 9 (x velocity), and 10 (y velocity). All of these images were generated using ParaView [41] from the computations with a reduced basis tolerance of 0.01 and ANOVA tolerance of 10^{-6} . Each figure consists of six subfigures arranged in a 2x3 grid. The columns correspond to the problems *iso*, *aniso1*, and *aniso2*, in that order. The first row depicts the mean error and the second row depicts the variance error.

The pressure errors are depicted in Fig. 7. The errors exhibit different behavior among the three problems. For the *iso* problem, the errors appear to be evenly distributed

Table 5 L^2 -norm of moment errors for reduced basis ANOVA estimates of the *iso* problem

ε^{RB}	ε^{A}			ε^{A}								
	10^{-4}	10^{-5}	10^{-6}	10^{-4}	10^{-5}	10^{-6}						
0.1	4.89×10^{-4}	4.50×10^{-4}	3.27×10^{-4}	2.42×10^{-4}	2.28×10^{-4}	2.43×10^{-4}						
	3.23×10^{-4}	2.10×10^{-4}	3.78×10^{-5}	3.65×10^{-5}	3.66×10^{-5}	3.88×10^{-5}						
	1.11×10^{-4}	9.82×10^{-5}	6.63×10^{-6}	3.61×10^{-6}	4.34×10^{-6}	7.29×10^{-6}						
0.01	3.41×10^{-1}	3.22×10^{-1}	2.41×10^{-2}	1.97×10^{-2}	1.97×10^{-2}	2.41×10^{-2}						
	2.03×10^{-1}	9.42×10^{-2}	2.71×10^{-3}	2.31×10^{-3}	2.31×10^{-3}	2.73×10^{-3}						
	5.49×10^{-3}	5.07×10^{-3}	7.11×10^{-4}	9.57×10^{-5}	1.01×10^{-4}	7.11×10^{-4}						
0.001	mean			variance								
combined												
pressure												
velocity												

Table 6 L^2 -norm of moment errors for reduced basis ANOVA estimates of the *aniso1* problem

		ε^A			ε^A			ε^A		
		10 ⁻⁴	10 ⁻⁵	10 ⁻⁶	10 ⁻⁴	10 ⁻⁵	10 ⁻⁶	10 ⁻⁴	10 ⁻⁵	10 ⁻⁶
ε_{RB}	1	4.49 × 10 ⁻⁴	3.95 × 10 ⁻⁴	3.95 × 10 ⁻⁴	7.70 × 10 ⁻⁵	4.65 × 10 ⁻⁵	4.66 × 10 ⁻⁵	7.74 × 10 ⁻⁵	4.70 × 10 ⁻⁵	4.71 × 10 ⁻⁵
	0.1	2.38 × 10 ⁻⁴	2.45 × 10 ⁻⁴	2.46 × 10 ⁻⁴	9.48 × 10 ⁻⁶	1.16 × 10 ⁻⁵	1.25 × 10 ⁻⁵	1.03 × 10 ⁻⁵	1.24 × 10 ⁻⁵	1.32 × 10 ⁻⁵
ε_{RB}	0.01	9.15 × 10 ⁻⁵	7.80 × 10 ⁻⁵	7.76 × 10 ⁻⁵	3.26 × 10 ⁻⁶	2.48 × 10 ⁻⁶	2.46 × 10 ⁻⁶	3.62 × 10 ⁻⁶	2.82 × 10 ⁻⁶	2.80 × 10 ⁻⁶
	1	9.35 × 10 ⁻²	6.82 × 10 ⁻²	6.82 × 10 ⁻²	9.51 × 10 ⁻³	8.28 × 10 ⁻³	8.22 × 10 ⁻³	9.53 × 10 ⁻³	8.29 × 10 ⁻³	8.23 × 10 ⁻³
ε_{RB}	0.1	1.06 × 10 ⁻²	1.03 × 10 ⁻²	1.04 × 10 ⁻²	1.27 × 10 ⁻³	6.07 × 10 ⁻⁴	6.04 × 10 ⁻⁴	1.27 × 10 ⁻³	6.10 × 10 ⁻⁴	6.07 × 10 ⁻⁴
	0.01	2.74 × 10 ⁻³	2.02 × 10 ⁻³	2.00 × 10 ⁻³	9.99 × 10 ⁻⁴	7.72 × 10 ⁻⁵	7.72 × 10 ⁻⁵	9.99 × 10 ⁻⁴	7.81 × 10 ⁻⁵	7.81 × 10 ⁻⁵
velocity		pressure			combined					

Table 7 L^2 -norm of moment errors for reduced basis ANOVA estimates of the *aniso2* problem

		ε^A			ε^A			ε^A		
		10 ⁻⁴	10 ⁻⁵	10 ⁻⁶	10 ⁻⁴	10 ⁻⁵	10 ⁻⁶	10 ⁻⁴	10 ⁻⁵	10 ⁻⁶
ε_{RB}	1	1.18 × 10 ⁻⁴	1.06 × 10 ⁻⁴	1.06 × 10 ⁻⁴	1.76 × 10 ⁻⁴	1.57 × 10 ⁻⁴	1.57 × 10 ⁻⁴	1.70 × 10 ⁻⁴	1.51 × 10 ⁻⁴	1.51 × 10 ⁻⁴
	0.1	3.79 × 10 ⁻⁵	4.24 × 10 ⁻⁵	4.24 × 10 ⁻⁵	2.47 × 10 ⁻⁵	2.42 × 10 ⁻⁵	2.42 × 10 ⁻⁵	2.67 × 10 ⁻⁵	2.71 × 10 ⁻⁵	2.71 × 10 ⁻⁵
ε_{RB}	0.01	2.51 × 10 ⁻⁵	1.48 × 10 ⁻⁵	1.49 × 10 ⁻⁵	1.73 × 10 ⁻⁵	9.29 × 10 ⁻⁶	7.56 × 10 ⁻⁶	1.85 × 10 ⁻⁵	1.01 × 10 ⁻⁵	8.82 × 10 ⁻⁶
	1	6.19 × 10 ⁻²	4.94 × 10 ⁻²	4.94 × 10 ⁻²	3.33 × 10 ⁻²	2.33 × 10 ⁻²	2.34 × 10 ⁻²	3.34 × 10 ⁻²	2.35 × 10 ⁻²	2.35 × 10 ⁻²
ε_{RB}	0.1	3.82 × 10 ⁻³	4.72 × 10 ⁻³	4.73 × 10 ⁻³	1.62 × 10 ⁻³	1.81 × 10 ⁻³	1.83 × 10 ⁻³	1.63 × 10 ⁻³	1.83 × 10 ⁻³	1.84 × 10 ⁻³
	0.01	2.24 × 10 ⁻³	1.35 × 10 ⁻³	1.58 × 10 ⁻³	8.36 × 10 ⁻⁴	3.58 × 10 ⁻⁴	4.03 × 10 ⁻⁴	8.44 × 10 ⁻⁴	3.65 × 10 ⁻⁴	4.12 × 10 ⁻⁴
velocity		pressure			combined					




Fig. 7 Moment errors in pressure for each of the three problems

throughout the domain although with clear increases at the corners of subdomains. However, for the *aniso1* (resp., *aniso2*), problem, we observe vertical (resp., horizontal) bands. Recall that the chosen permeabilities for the anisotropic problems favor either vertical or horizontal flow. The bands appear to reflect these favored permeabilities, as higher errors might be expected where values are greater. As with the *iso* case, interfaces and corners between subdomains are emphasized. The heightened errors at these interfaces are likely due to higher order effects not captured in the ANOVA terms selected for the expansion.

The velocity errors are depicted in Figs. 8, 9, and 10. As with the pressures errors, the errors in all cases are heightened along subdomain interfaces, as expected due to the exclusion of higher order ANOVA terms. Furthermore, the favored permeabilities for the anisotropic problems result in larger errors in the velocity directions which are favored, again likely due to the values being larger.

While the above analysis demonstrates that a smaller reduced basis tolerance greatly improves the accuracy of the moment computations, it comes with a cost of a larger number of high-fidelity solves and a larger reduced basis. Table 8 summarizes the number of high-fidelity solves for each problem and each choice of tolerances. For comparison, the number of collocation points is provided in parentheses. For all three problems, it appears that the number of high-fidelity solves almost doubles as the reduced basis tolerance is reduced by a factor of 10. However, the number of high-fidelity solves remains fairly constant even as the ANOVA tolerance increases. This suggests that a small reduced basis is able to fairly accurately approximate the space of high-fidelity solves.

The decrease in the a posteriori error estimates over each iteration of training is depicted in Fig. 11. The velocity error estimate is depicted as the solid blue line, the pressure error estimate as the red dotted line, and the combined




Fig. 8 Moment errors in velocity magnitude for each of the three problems

error estimate as the orange dashed line. Each error estimate depicted is normalized by the norm of the reduced basis solution to produce a relative error estimate. The figure displays, at each iteration, the largest relative error estimate of the given type. Since the largest error estimates for each type are not guaranteed to be taken from the same sample at any given iteration, and each of these error estimates are normalized by the reduced basis solution at that sample, there is no guaranteed ordering of the values (i.e., the combined error is not necessarily greater than the velocity or pressure error).

For the *iso* problem, all three errors appear to be roughly equal. For the *aniso1* problem, however, the velocity error is the largest with the combined error being in between velocity and pressure. Recall that the *aniso1* problem favors vertical flow, but the boundary conditions impose horizontal flow. This contention between vertical and horizontal flow perhaps explains the increased difficulty of the reduced

basis in approximating the velocity space. On the other hand, the pressure error is the greatest for the *aniso2* problem with the velocity and combined errors being roughly equal. For this problem, horizontal flow is favored so that there is no contention between the flow specified by the boundary conditions and the permeabilities, allowing easier approximation of velocity.

The error estimates are displayed on a log scale, suggesting linear convergence. The black vertical line marks the iteration at which training terminated after the first ANOVA level. The increase in the error estimates at the start of the second ANOVA level is due to an increase in the number of collocation points used during training. In both the *iso* and *aniso1* problems, the increase in error between levels is small, suggesting that the reduced basis formed at the end of the first level generalized fairly well to the new collocation points. The *aniso2* problem has a larger increase in error, but there was also a




Fig. 9 Moment errors in the x velocity for each of the three problems

larger decrease in error towards the end of the level 1 training. This suggests that the level 1 reduced basis did not generalize well to the new collocation points. In all cases, however, the error appears to decrease during the second level at roughly the same rate as during the first level, despite the greater increase in the number of training samples.

The a posteriori error estimates only compute an upper bound on the actual error. It would be helpful to understand how tight these error estimates actually are. To study this, we selected a random subset of 1000 ANOVA collocation points from each problem using tolerances $\varepsilon^{\text{RB}} = 0.01$ and $\varepsilon^{\text{A}} = 10^{-6}$ and computed high-fidelity solutions at each selected point. Using the associated reduced basis for each problem, we also computed the reduced basis approximation. From this, we obtained exact errors for each collocation point and computed ratios of the a posteriori error estimate against the exact error. This ratio is called the *effectivity* of the error estimator. These values are presented in Fig. 12. Velocity ratios are shown in blue upward facing

triangles, pressure ratios in red downward facing triangles, and combined ratios in orange circles. Since the error estimates form upper bounds, we expect these ratios to be greater than 1. However, if the bound is tight, we expect the ratios to be close to 1. From the figure, it is apparent that all ratios are greater than 1, suggesting our error estimates are true upper bounds. However, the ratio is much greater than 1, suggesting that the upper bounds are not as tight as we would like them to be for an ideal setup. It is well known that the tightness of the error estimate is related to the condition number [23, Section 3.6.2]. As demonstrated in Table 2, the Stokes-Brinkman systems are highly ill-conditioned, resulting in loose upper bounds with the error estimates. Thus, the effectivities of Fig. 12 are expected. Nonetheless, despite the lack of sharpness in the upper bounds, our experiments provide numerical evidence that combining reduced basis methods with anchored ANOVA yields substantial benefits to estimating the moments of stochastic Stokes-Brinkman problems.




Fig. 10 Moment errors in the y velocity for each of the three problems

Table 8 Comparison of number of high-fidelity solves used to form the reduced basis and number of ANOVA collocation points (in parentheses) for the three problems with varying tolerances

ϵ^A		10^{-4}	10^{-5}	10^{-6}	
ϵ^{RB}	1	28 (1573)	33 (27701)	33 (43541)	iso
	0.1	59 (4021)	61 (33605)	62 (50885)	
	0.01	110 (7285)	122 (43541)	123 (52165)	
ϵ^{RB}	1	29 (1537)	31 (19889)	31 (37825)	aniso1
	0.1	50 (2737)	52 (27665)	52 (40049)	
	0.01	90 (3985)	95 (36737)	96 (41185)	
ϵ^{RB}	1	34 (2737)	37 (35665)	37 (41185)	aniso2
	0.1	73 (2737)	77 (34609)	77 (41185)	
	0.01	91 (3329)	120 (34609)	121 (41185)	

7 Conclusion

We demonstrated that the use of the truncated ANOVA decomposition is effective in reducing the number of collocation points needed for accurate approximation of the statistical moments of several Stokes-Brinkman problems with stochastic permeability. Additionally, we showed that reduced basis methods yield significant savings by reducing the number of high-fidelity solves required to compute these moments. The reduced basis methods rely upon accurate and efficient a posteriori error estimates. We present such estimates based upon the Brezzi stability theory. While these error estimates allowed us to construct small reduced bases, they were not as sharp as desired due to the high-fidelity systems being highly ill-conditioned. Nonetheless, excellent reduced basis approximations were obtained.




Fig. 11 Maximum values of error estimates for each training iteration using $\varepsilon^A = 10^{-6}$ and $\varepsilon^{RB} = 0.01$




Fig. 12 Ratio of the error estimates (34)–(38) against the exact errors from high-fidelity solves for each of the three problems at 1000 randomly sampled ANOVA collocation points

Acknowledgements We would like to thank Prof. Howard C. Elman for sharing his notes regarding inf-sup stability for reduced basis methods for saddle-point problems.

Declarations

Conflict of Interests The authors declare that they have no conflict of interest.

References

1. Arbogast, T., Brunson, D.S.: A computational method for approximating a Darcy-Stokes system governing a vuggy porous medium. *Comput. Geosci.* **11**(3), 207–218 (2007). <https://doi.org/10.1007/s10596-007-9043-0>
2. Beavers, G.S., Joseph, D.D.: Boundary conditions at a naturally permeable wall. *J. Fluid Mech* **30**, 197–207 (1967)
3. Urquiza, J.M., N'Dri, D., Garon, A., Delfour, M.C.: Coupling Stokes and Darcy equations. *Appl. Numer. Math.* **58**(5), 525–538 (2008). <https://doi.org/10.1016/j.apnum.2006.12.006>, <http://www.sciencedirect.com/science/article/pii/S016892740600225X>
4. Darcy, H.: Les fontaines publiques de la ville de Dijon. Dalmont, Paris (1856)
5. Whitaker, S.: Flow in porous media I: A theoretical derivation of Darcy's law. *Transp. Porous Media* **1**(1), 3–25 (1986). <https://doi.org/10.1007/BF01036523>
6. Bear, J.: Dynamics of fluids in porous media. Elsevier (1972)
7. Brinkman, H.C.: A calculation of the viscous force exerted by a flowing fluid on a dense swarm of particles. *Appl. Sci. Res.* **A1**, 27–34 (1948)
8. Gulbransen, A.F., Hauge, V.L., Lie, K.-A.: A multiscale mixed finite-element method for vuggy and naturally-fractured reservoirs. *SPE J.* **15**(2), 395–403 (2010)
9. Laptev, V.: Numerical solution of coupled flow in plain and porous media. Ph.D. Thesis, Technical University of Kaiserslautern (2003)
10. Popov, P., Bi, L., Efendiev, Y., Ewing, R.E., Qin, G., Li, J., Ren, Y.: Multi-physics and multi-scale methods for modeling fluid flow through naturally-fractured vuggy carbonate reservoirs. In: Proceedings of the 15th SPE Middle East Oil & Gas Show and Conference, Bahrain. (Paper SPE 105378) (2007)
11. Ghanem, R.G., Spanos, P.D.: Stochastic finite elements: A spectral approach. Springer, Berlin (1991)
12. Papoulis, A.: Probability, random variables, and stochastic processes. Communications and signal processing. McGraw-Hill. <https://books.google.com/books?id=4IwQAQAAIAAJ> (1991)
13. Xiu, D., Hesthaven, J.: High-order collocation methods for differential equations with random inputs. *SIAM J. Sci. Comput.* **27**, 1118–1139 (2005). <https://doi.org/10.1137/040615201>
14. Nobile, F., Tempone, R., Webster, C.: A sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM J. Numer. Anal.* **46**, 2309–2345 (2008). <https://doi.org/10.1137/060663660>
15. Cao, Y., Chen, Z., Gunzburger, M.: ANOVA expansions and efficient sampling methods for parameter dependent nonlinear PDEs. *Int. J. Numer. Anal. Model.* (2009)
16. Fisher, R.A.: Statistical methods for research workers. Edinburgh Oliver & Boyd (1925)
17. Foo, J., Karniadakis, G.E.: Multi-element probabilistic collocation method in high dimensions. *J. Comput. Phys.* **229**(5), 1536–1557 (2010). <https://doi.org/10.1016/j.jcp.2009.10.043>
18. Ma, X., Zabaras, N.: An adaptive high-dimensional stochastic model representation technique for the solution of stochastic partial differential equations. *J. Comput. Phys.* **229**(10), 3884–3915 (2010). <https://doi.org/10.1016/j.jcp.2010.01.033>
19. Yang, X., Choi, M., Lin, G., Karniadakis, G.: Adaptive ANOVA decomposition of stochastic incompressible and compressible flows. *J. Comput. Physics* **231**, 1587–1614 (2012). <https://doi.org/10.1016/j.jcp.2011.10.028>
20. Boyaval, S., Bris, C., Lelièvre, T., Maday, Y., Nguyen, N., Patera, A.: Reduced basis techniques for stochastic problems. *Arch. Comput. Methods Eng.* **17**, 435–454 (2010). <https://doi.org/10.1007/s11831-010-9056-z>
21. Haasdonk, B., Ohlberger, M.: Reduced basis method for finite volume approximations of parametrized linear evolution equations. *ESAIM: Math. Modell. Numer. Anal.* **42**, 277–302 (2008). <https://doi.org/10.1051/m2an:2008001>
22. Hesthaven, J., Rozza, G., Stamm, B.: Certified reduced basis methods for parametrized partial differential equations (2016)
23. Quarteroni, A., Manzoni, A., Negri, F.: Reduced basis methods for partial differential equations: An introduction (2015)
24. Hesthaven, J.S., Zhang, S.: On the use of ANOVA expansions in reduced basis methods for parametric partial differential equations. *J. Sci. Comput.* **69**(1), 292–313 (2016). <https://doi.org/10.1007/s10915-016-0194-9>
25. Xu, H., Rahman, S.: A generalized dimension?reduction method for multidimensional integration in stochastic mechanics. *Int. J. Numer. Methods Eng.* **61**, 1992–2019 (2004). <https://doi.org/10.1002/nme.11>
26. Zang, H., Choi, M., Karniadakis, G.: Error estimates for the ANOVA method with polynomial chaos interpolation: Tensor product functions. *SIAM J. Sci. Comput.* **34**, 1165–1186 (2012). <https://doi.org/10.1137/100788859>
27. Zhang, Z., Choi, M., Karniadakis, G.E.: Anchor points matter in ANOVA decomposition. In: Hesthaven, J.S., Rønquist, E.M. (eds.) Spectral and High Order Methods for Partial Differential Equations, pp. 347–355. Springer, Berlin (2011)
28. Cho, H., Elman, H.: An adaptive reduced basis collocation method based on PCM ANOVA decomposition for anisotropic stochastic PDEs. *Int. J. Uncertain. Quantif.* **8**. <https://doi.org/10.1615/Int.J.UncertaintyQuantification.2018024436> (2017)
29. Liao, Q., Lin, G.: Reduced basis ANOVA methods for partial differential equations with high-dimensional random inputs. *J. Comput. Phys.* **317**. <https://doi.org/10.1016/j.jcp.2016.04.029> (2016)
30. Rozza, G., Veroy, K.: On the stability of reduced basis methods for Stokes equations in parametrized domains. *Comput. Methods Appl. Mech. Eng.* **196**(7), 1244–1260 (2007). <https://doi.org/10.1016/j.cma.2006.09.005>, <http://infoscience.epfl.ch/record/103016>
31. Gerner, A.-L., Veroy, K.: Certified reduced basis methods for parametrized saddle point problems. *SIAM J. Sci. Comput.* **34**(5), A2812–A2836 (2012). <https://doi.org/10.1137/110854084>
32. Boffi, D., Brezzi, F., Fortin, M.: Mixed finite element methods and applications. Springer Series in Computational Mathematics. Springer, Berlin. <https://books.google.com/books?id=mRhAAAAAQBAJ> (2013)
33. Brezzi, F., Fortin, M.: Mixed and hybrid finite element methods. Springer, Berlin (1991)
34. Elman, H.C., Silvester, D.J., Wathen, A.J.: Finite elements and fast iterative solvers: with applications in incompressible fluid dynamics, 2nd edn. Oxford University Press, New York (2014)
35. Tang, K., Congedo, P.M., Abgrall, R.: Sensitivity analysis using anchored ANOVA expansion and high-order moments computation. *Int. J. Numer. Methods Eng.* **102**, 1554–1584 (2015). <https://doi.org/10.1002/nme.4856>

36. Elman, H.C., Liao, Q.: Reduced basis collocation methods for partial differential equations with random coefficients. *SIAM/ASA J. Uncertain. Quant.* **1**, 192–217 (2013)
37. Huynh, D.B.P., Rozza, G., Sen, S., Patera, A.T.: A successive constraint linear optimization method for lower bounds of parametric coercivity and inf-sup stability constants. *C. R. Math. Acad. Sci. Paris* **345**(8), 473–478 (2007). <https://doi.org/10.1016/j.crma.2007.09.019>
38. Chen, Y., Hesthaven, J.S., Maday, Y., Rodríguez, J.: Improved successive constraint method based a posteriori error estimate for reduced basis approximation of 2D Maxwell's problem. *ESAIM: M2AN* **43**(6), 1099–1116 (2009). <https://doi.org/10.1051/m2an/2009037>
39. Knyazev, A.: Toward the optimal preconditioned eigensolver: Locally optimal block preconditioned conjugate gradient method. *SIAM J. Sci. Comput.* **23**, 517–541 (2001). <https://doi.org/10.1137/S1064827500366124>
40. Knyazev, A.: Locally optimal block preconditioned conjugate gradient. GitHub. <https://github.com/lobpcg/blopex> (2019)
41. Ayachit, U.: The paraview guide: A parallel visualization application. Kitware (2015)

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.