

MATHICSE Technical Report

Nr. 10.2015

April 2015



Subspace acceleration for large-scale parameter-dependent Hermitian eigenproblems

Petar Sirković, Daniel Kressner

Subspace acceleration for large-scale parameter-dependent Hermitian eigenproblems ^{*}

Petar Sirković[†]

Daniel Kressner[‡]

April 16, 2015

Abstract

This work is concerned with approximating the smallest eigenvalue of a parameter-dependent Hermitian matrix $A(\mu)$ for many parameter values $\mu \in \mathbb{R}^P$. The design of reliable and efficient algorithms for addressing this task is of importance in a variety of applications. Most notably, it plays a crucial role in estimating the error of reduced basis methods for parametrized partial differential equations. The current state-of-the-art approach, the so called Successive Constraint Method (SCM), addresses affine linear parameter dependencies by combining sampled Rayleigh quotients with linear programming techniques. In this work, we propose a subspace approach that additionally incorporates the sampled eigenvectors of $A(\mu)$ and implicitly exploits their smoothness properties. Like SCM, our approach results in rigorous lower and upper bounds for the smallest eigenvalues on D . Theoretical and experimental evidence is given to demonstrate that our approach represents a significant improvement over SCM in the sense that the bounds are often much tighter, at negligible additional cost.

Keywords. parameter-dependent eigenvalue problem, Hermitian matrix, subspace acceleration, Successive Constraint Method, quadratic residual bound

1 Introduction

Let $A : D \rightarrow \mathbb{C}^{N \times N}$ be a matrix-valued function on a compact subset $D \subset \mathbb{R}^P$ such that $A(\mu)$ is Hermitian for every $\mu \in D$. We aim at approximating the smallest eigenvalue,

$$\lambda_{\min}(A(\mu)), \quad \mu \in D, \quad (1)$$

of $A(\mu)$ for *many* different values of $\mu \in D$. We consider a large-scale setting, where applying a standard eigensolver, such as the Lanczos method [2], is computationally feasible for a few values of μ but would become too expensive for many (e.g., thousand) parameter values. Guiding our developments, an important application of (1) consists of estimating the coercivity constant for parametrized elliptic partial differential equations (PDEs), see, e.g., [34]. In turn, these estimates can be used to construct

^{*}Supported by the SNF research module *A Reduced Basis Approach to Large-Scale Pseudospectra Computations* within the SNF ProDoc *Efficient Numerical Methods for Partial Differential Equations*.

[†]ANCHP, MATHICSE, EPF Lausanne, Switzerland. petar.sirkovic@epfl.ch

[‡]ANCHP, MATHICSE, EPF Lausanne, Switzerland. daniel.kressner@epfl.ch

reliable a posteriori error estimates in the reduced basis method (RBM) for solving such PDEs. For more general PDEs, the coercivity constant needs to be replaced by the inf-sup constant, which – after discretization – corresponds to the smallest singular value of a general nonsymmetric matrix $A(\mu)$ or, equivalently, to the smallest eigenvalue of the Hermitian matrix $A(\mu)^*A(\mu)$. Other applications that require the solution of such parameter-dependent eigenvalue and singular value problems include the computation of pseudospectra [37, Part IX], the method of particular solutions [4], and the eigenvalue analysis of waveguides [6]. The related problem of optimizing the smallest eigenvalue(s) of a parameter-dependent Hermitian matrix appears in a large variety of applications: One-parameter optimization problems play a critical role in the design of numerical methods [33] and robust control [21]; multi-parameter optimization problems arise from semidefinite programming [9] and graph partitioning [16, 7].

Without any further assumptions on the dependence of $A(\mu)$ on μ , the solution of (1) is computationally intractable, especially when P is large. An assumption commonly found in RBM is that $A(\mu)$ admits an affine linear decomposition with respect to μ .

Assumption 1.1 (Affine linear decomposition). *Given $Q \in \mathbb{N}$, the Hermitian matrix $A(\mu)$ admits a decomposition of the form*

$$A(\mu) = \theta_1(\mu)A_1 + \dots + \theta_Q(\mu)A_Q, \quad \forall \mu \in D, \quad (2)$$

for Hermitian matrices $A_1, \dots, A_Q \in \mathbb{C}^{N \times N}$ and functions $\theta_1, \dots, \theta_Q : D \mapsto \mathbb{R}$.

Assumption 1.1 holds with small Q for a number of important applications, including PDEs with parametrized coefficients on disjoint subdomains [34]. Even when $A(\mu)$ does not satisfy this assumption, it may still be possible to approximate it very well by a short affine linear decomposition using, e.g., the Empirical Interpolation Method [3].

One of the simplest approaches to address (1) is to use Gershgorin’s theorem [14] for estimating the smallest eigenvalue, but the accuracy of the resulting estimate is usually insufficient and limits the scope of applications severely. Within the context of RBM, a number of approaches have been developed that go beyond this simple estimate by making use of Assumption 1.1. For example, eigenvalue perturbation analysis can be used to locally approximate the smallest eigenvalues [28, 38]. The Successive Constraint Method (SCM; see [12]) is currently the most commonly used approach within RBM, probably due to its generality and relative simplicity. Variants of SCM for computing smallest singular values can be found in [35, 11], while an extension of SCM to non-linear problems and alternative heuristic strategies have been proposed in [24].

If A depends analytically on μ then the smallest eigenvalue inherits this property if $\lambda_{\min}(A(\mu))$ remains simple [15]. As shown in [1], the analyticity can be used to approximate $\lambda_{\min}(A(\mu))$ very well by high-order Legendre polynomials (for $P = 1$) or sparse tensor products of Legendre polynomials (for $P > 1$ if D is a hypercube). Requiring $\lambda_{\min}(A(\mu))$ to stay simple on the whole of D is, however, a rather strong condition. In general, there are eigenvalue crossings at which $\lambda_{\min}(A(\mu))$ is Lipschitz continuous only; see [27] for a recently proposed eigenvalue optimization method that takes this piecewise regularity of $\lambda_{\min}(A(\mu))$ into account. For larger P , keeping track of eigenvalue crossings explicitly appears to be a rather daunting task and we therefore aim at a method for solving (1) that benefits only implicitly from piecewise regularity.

The approach proposed in this paper can be summarized as follows. Given J parameter samples $\mu_1, \mu_2, \dots, \mu_J$, we consider the subspace \mathcal{V} containing eigenvectors belonging

to one or several smallest eigenvalues of $A(\mu_i)$ for $i = 1, \dots, J$. The smallest Ritz value of $A(\mu)$ with respect to \mathcal{V} immediately yields an upper bound for $\lambda_{\min}(A(\mu))$. A lower bound is obtained by combining this upper bound with a perturbation argument. To apply such an argument requires, however, knowledge on the involved eigenvalue gap. We show that this gap can be estimated by adapting the linear programming approach used in SCM. The difference between upper and lower bounds constitutes the error estimate that drives the greedy strategy for selecting the next parameter sample μ_{J+1} . The whole procedure is stopped once the error estimate is uniformly small on D or, rather, on a surrogate of D .

As we will see in the numerical experiments, our subspace approach accelerates convergence significantly compared to SCM. Subspace approaches based on additional conditions on the parameter dependencies have been proposed in [23, 31]. In the context of eigenvalue optimization problems, subspace acceleration has been discussed in [5, 19].

The rest of this paper is organized as follows. In Section 2, we first give a brief overview of SCM. We then discuss its interpolation properties and point out a limitation on the quality of the lower bounds that can possibly be attained when solely using the information taken into account by SCM. In Section 3, we present our novel subspace-accelerated approach for solving (1). Furthermore, we show that the new approach has better interpolation properties than SCM. Motivated by the fast convergence of the upper bounds in the novel approach, we also introduce residual-based lower bounds which are less reliable but sometimes converge much faster. In Sections 4 and 5, we present numerical experiments and discuss the application of our algorithm to the computation of coercivity and inf-sup constants.

2 Successive Constraint Method

In the following, we recall the Successive Constraint Method (SCM) from [12] and derive some new properties. The basic idea of SCM is to construct reduced-order models for (1) that allow the efficient evaluation of lower and upper bounds for $\lambda_{\min}(A(\mu))$. Being a consequence of Assumption 1.1, the following characterization of the smallest eigenvalue is central to SCM:

$$\begin{aligned} \lambda_{\min}(A(\mu)) &= \min_{\substack{u \in \mathbb{C}^N \\ u \neq 0}} \frac{u^* A(\mu) u}{u^* u} = \min_{\substack{u \in \mathbb{C}^N \\ u \neq 0}} \sum_{q=1}^Q \theta_q(\mu) \frac{u^* A_q u}{u^* u} \\ &= \min_{\substack{u \in \mathbb{C}^N \\ u \neq 0}} \theta(\mu)^T R(u) = \min_{y \in \mathcal{Y}} \theta(\mu)^T y, \end{aligned} \quad (3)$$

where we have defined the vector-valued functions $\theta : D \rightarrow \mathbb{R}^Q$, $R : \mathbb{C}^N \setminus \{0\} \rightarrow \mathbb{R}^Q$ as

$$\theta(\mu) := [\theta_1(\mu), \dots, \theta_Q(\mu)]^T, \quad R(u) := \left[\frac{u^* A_1 u}{u^* u}, \dots, \frac{u^* A_Q u}{u^* u} \right]^T, \quad (4)$$

and set $\mathcal{Y} := \text{im}(R)$. It follows from (3) that the computation of $\lambda_{\min}(A(\mu))$ is equivalent to optimizing a linear functional over \mathcal{Y} . The constraint set \mathcal{Y} is called the *joint numerical range* of A_1, \dots, A_Q , which is generally not convex; see [8]. Thus standard optimization techniques cannot be used to reliably solve (3). In SCM, the set \mathcal{Y} is approximated from above and from below by convex polyhedra. In turn, this allows for the use of linear programming (LP) techniques to yield lower and upper bounds.

2.1 Basic idea of SCM

Given J parameter values $C_J = \{\mu_1, \dots, \mu_J\} \subset D$, let us suppose we have computed the corresponding eigenpairs $(\lambda_1, v_1), \dots, (\lambda_J, v_J)$, that is, λ_i is the smallest eigenvalue of $A(\mu_i)$ with eigenvector $v_i \in \mathbb{C}^N$. We now describe how SCM uses this information to approximate the set \mathcal{Y} defined above.

Clearly,

$$\mathcal{Y}_{\text{UB}}(C_J) := \{R(v_i) : i = 1, \dots, J\} \quad (5)$$

is a subset of \mathcal{Y} . Optimizing (4) over $\mathcal{Y}_{\text{UB}}(C_J)$ instead of \mathcal{Y} thus yields an *upper bound* for $\lambda_{\min}(A(\mu))$. Note that this is equivalent to optimizing over the convex hull of $\mathcal{Y}_{\text{UB}}(C_J)$, since a solution of the LP can always be attained at a vertex of the convex polyhedron.

To get a lower bound, we first define the bounding box

$$\mathcal{B} = [\lambda_{\min}(A_1), \lambda_{\max}(A_1)] \times \dots \times [\lambda_{\min}(A_Q), \lambda_{\max}(A_Q)] \subseteq \mathbb{R}^Q. \quad (6)$$

By the minimax characterization of eigenvalues we have $\mathcal{Y} \subset \mathcal{B}$, but this approximation is often too crude and we will instead work with

$$\mathcal{Y}_{\text{LB}}(C_J) := \{y \in \mathcal{B} : \theta(\mu_i)^T y \geq \lambda_i, i = 1, \dots, J\}.$$

The property $\mathcal{Y} \subset \mathcal{Y}_{\text{LB}}(C_J)$ follows from the fact that every $y = R(u_y) \in \mathcal{Y}$ satisfies $\theta(\mu_i)^T y = u_y^* A(\mu_i) u_y / u_y^* u_y \geq \min_u u^* A(\mu_i) u / u^* u = \lambda_i$. The minimax characterization also implies that the convex polyhedron $\mathcal{Y}_{\text{LB}}(C_J)$ is tangential to \mathcal{Y} .

With the sets defined above, we let

$$\lambda_{\text{UB}}(\mu; C_J) := \min_{y \in \mathcal{Y}_{\text{UB}}(C_J)} \theta(\mu)^T y, \quad \lambda_{\text{LB}}(\mu; C_J) := \min_{y \in \mathcal{Y}_{\text{LB}}(C_J)} \theta(\mu)^T y. \quad (7)$$

Since $\mathcal{Y}_{\text{UB}}(C_J) \subseteq \mathcal{Y} \subseteq \mathcal{Y}_{\text{LB}}(C_J)$, it follows that

$$\lambda_{\text{LB}}(\mu; C_J) \leq \lambda_{\min}(A(\mu)) \leq \lambda_{\text{UB}}(\mu; C_J)$$

for every $\mu \in D$. While the evaluation of $\lambda_{\text{UB}}(\mu; C_J)$ is trivial, the evaluation of $\lambda_{\text{LB}}(\mu; C_J)$ requires the solution of an LP; see Figure 1 for an illustration.

2.2 Error estimate and sampling strategy

Assessing the quality of the bounds (7) on the entire, usually continuous parameter domain D is, in general, an infeasible task. A common strategy in SCM, we substitute D by a training set $\Xi \subset D$ that contains finitely many (usually, a few thousand) parameter samples. We then measure the quality of the bounds by estimating the largest relative difference:

$$\max_{\mu \in \Xi} \frac{\lambda_{\text{UB}}(\mu; C_J) - \lambda_{\text{LB}}(\mu; C_J)}{|\lambda_{\text{UB}}(\mu; C_J)|}. \quad (8)$$

If (8) is not sufficiently small, SCM enlarges C_J by a parameter that attains the maximum in (8) and recomputes the bounds (7). The resulting greedy sampling strategy is summarized in Algorithm 1.




Figure 1: Illustration of the LP defining the lower bound $\lambda_{\text{LB}}(\mu; C_J)$ for $Q = 2$ and $J = 2$.

Algorithm 1 Successive Constraint Method

Input: Training set Ξ , affine linear decomposition such that $A(\mu) = \theta_1(\mu)A_1 + \dots + \theta_Q(\mu)A_Q$ is Hermitian for every $\mu \in \Xi$. Relative error tolerance ε_{SCM} .

Output: Set $C_J \subset \Xi$ with corresponding eigenpairs (λ_i, v_i) , such that

$$\frac{\lambda_{\text{UB}}(\mu; C_J) - \lambda_{\text{LB}}(\mu; C_J)}{|\lambda_{\text{UB}}(\mu; C_J)|} < \varepsilon_{\text{SCM}} \text{ for every } \mu \in \Xi.$$

- 1: compute $\lambda_{\min}(A_q), \lambda_{\max}(A_q)$ for $q = 1, \dots, Q$, defining \mathcal{B} according to (6)
 - 2: $J = 0, C_0 = \emptyset$
 - 3: **while** $\max_{\mu \in \Xi} \frac{\lambda_{\text{UB}}(\mu; C_J) - \lambda_{\text{LB}}(\mu; C_J)}{|\lambda_{\text{UB}}(\mu; C_J)|} > \varepsilon_{\text{SCM}}$ **do**
 - 4: $\mu_{J+1} \leftarrow \arg \max_{\mu \in \Xi} \frac{\lambda_{\text{UB}}(\mu; C_J) - \lambda_{\text{LB}}(\mu; C_J)}{|\lambda_{\text{UB}}(\mu; C_J)|}$
 - 5: $C_{J+1} \leftarrow C_J \cup \mu_{J+1}$
 - 6: recompute $\lambda_{\text{UB}}(\mu; C_{J+1})$ and $\lambda_{\text{LB}}(\mu; C_{J+1})$ according to (7)
 - 7: $J \leftarrow J + 1$
 - 8: **end while**
-

2.3 Computational complexity

Let us briefly summarize the computations performed by SCM. The bounding box \mathcal{B} for \mathcal{Y} needs to be determined initially by computing the smallest and the largest eigenvalues of A_1, \dots, A_Q . Since each iteration requires the computation of the smallest eigenpair (λ_i, v_i) of $A(\mu_i)$, this amounts to solving $2Q + J$ eigenproblems of size $N \times N$ in total. Verifying the accuracy of the current approximation on Ξ and selecting the next parameter sample requires computing $\lambda_{\text{UB}}(\mu; C_J)$ and $\lambda_{\text{LB}}(\mu; C_J)$ for all $\mu \in \Xi$. In total, this amounts to solving $J|\Xi|$ LP problems with Q variables and at most $2Q + J$ constraints.

2.4 Interpolation results

As also discussed in [12], it is immediate to see that the bounds produced by SCM coincide with $\lambda_{\min}(A(\mu))$ for all $\mu \in C_J$. The following theorem shows that the upper bounds also interpolate the derivatives of $\lambda_{\min}(A(\mu))$ on C_J .

Theorem 2.1. Let $C_J \subset D$ be finite and consider the upper bound $\lambda_{\text{UB}}(\mu; C_J)$ defined in (7). Given $\mu_i \in C_J$ in the interior of D , assume that $\theta_1, \dots, \theta_Q : D \rightarrow \mathbb{R}$ are differentiable at μ_i and that $\lambda_i = \lambda_{\min}(A(\mu_i))$ is a simple eigenvalue of $A(\mu_i)$. Then

$$\nabla \lambda_{\text{UB}}(\mu_i; C_J) = \nabla \lambda_{\min}(A(\mu_i)),$$

with the gradient ∇ with respect to μ .

Proof. Let v_i be an eigenvector associated with λ_i such that $\|v_i\|_2 = 1$ and set $y_i := R(v_i) \in \mathcal{Y}_{\text{UB}}(C_J)$. By the definition (7), the relation

$$\lambda_{\text{UB}}(\mu; C_J) = \min_{y \in \mathcal{Y}_{\text{UB}}(C_J)} \theta(\mu)^T y = \theta(\mu)^T y_i \quad (9)$$

holds for $\mu = \mu_i$. The simplicity of λ_i implies that y_i is the unique minimizer. Combined with the facts that $\mathcal{Y}_{\text{UB}}(C_J)$ is a discrete set, μ_i is an interior point, and $\theta(\mu)$ is continuous at μ_i , this implies that (9) also holds for all μ in a neighbourhood $\Omega \subset D$ around μ_i . Consequently,

$$\frac{\partial \lambda_{\text{UB}}}{\partial \mu^{(p)}}(\mu_i; C_J) = \frac{\partial \theta}{\partial \mu^{(p)}}(\mu_i)^T y_i,$$

where $\mu^{(p)}$ denotes the p th entry of μ for $p = 1, \dots, P$.

On the other hand, the well-known expression for the derivative of a simple eigenvalue [10] gives

$$\begin{aligned} \frac{\partial \lambda_{\min}}{\partial \mu^{(p)}}(A(\mu_i)) &= v_i^* \frac{\partial A}{\partial \mu^{(p)}}(\mu_i) v_i = v_i^* \left(\sum_{q=1}^Q \frac{\partial \theta_q}{\partial \mu^{(p)}}(\mu_i) A_q \right) v_i \\ &= \sum_{q=1}^Q \frac{\partial \theta_q}{\partial \mu^{(p)}}(\mu_i) v_i^* A_q v_i = \frac{\partial \theta}{\partial \mu^{(p)}}(\mu_i)^T y_i, \end{aligned}$$

which completes the proof. \square

As the following example shows, the result of Theorem 2.1 does not hold for the lower bounds produced by SCM.

Example 2.2. For $\mu \in D := [0, \pi]$, let

$$A(\mu) = \cos(\mu)A_1 + \sin(\mu)A_2 = \cos(\mu) \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} + \sin(\mu) \begin{bmatrix} 0 & -1 \\ -1 & 0 \end{bmatrix}.$$

It can be shown that \mathcal{Y} , the joint numerical range of A_1 and A_2 , equals the unit circle around 0. Consider the sample set $C_3 = \{\mu_1, \mu_2, \mu_3\} = \{0, \frac{\pi}{2}, \pi\}$, with $\lambda_{\min}(A(\mu_1)) = \lambda_{\min}(A(\mu_2)) = \lambda_{\min}(A(\mu_3)) = -1$. The resulting lower bound set $\mathcal{Y}_{\text{LB}}(C_3)$ is the half-infinite box shown in Figure 2. When minimizing $\theta(\mu)^T y$ for $y \in \mathcal{Y}_{\text{LB}}(C_3)$, the minimum is attained at the vertex $(-1, -1)$ for $\mu \in [\pi/4, \pi/2]$ and at the vertex $(1, -1)$ for $\mu \in [\pi/2, 3\pi/4]$. Hence,

$$\lambda_{\text{LB}}(\mu) = \begin{cases} -\cos \mu - \sin \mu, & \text{for } \mu \in [\pi/4, \pi/2], \\ \cos \mu - \sin \mu, & \text{for } \mu \in [\pi/2, 3\pi/4], \end{cases}$$




Figure 2: Joint numerical range \mathcal{Y} (red circle) and lower bound set $\mathcal{Y}_{\text{LB}}(C_J)$ (yellow area) for the setting described in Example 2.2.

yielding the following one-sided derivatives at $\mu = \pi/2$:

$$\begin{aligned}\lambda'_{\text{LB}}((\pi/2)^-) &= \sin(\pi/2) - \cos(\pi/2) = 1, \\ \lambda'_{\text{LB}}((\pi/2)^+) &= -\sin(\pi/2) - \cos(\pi/2) = -1.\end{aligned}$$

In contrast, the exact eigenvalue is differentiable at $\pi/2$. Moreover, $\lambda'_{\min}(A(\pi/2)) = 0$ is different from both one-sided derivatives of $\lambda_{\text{LB}}(\mu)$.

Theorem 2.1 and Example 2.2 indicate that the lower bounds produced by SCM are asymptotically less accurate than the upper bounds. Indeed, this has been observed numerically [11], implying a need to find more accurate lower bounds. However, the following theorem indicates that such an improvement is not possible without taking additional information on $A(\mu)$ into account.

Theorem 2.3. *Let $C_J = \{\mu_1, \dots, \mu_J\} \subseteq D$ and consider the lower bounds $\lambda_{\text{LB}}(\mu; C_J)$ defined in (7) for a Hermitian matrix function $A(\mu)$ in affine linear decomposition (2). Let $\tilde{\mu} \in D$. If $J < N$ then there exist matrices $\bar{A}_1, \dots, \bar{A}_Q \in \mathbb{C}^{N \times N}$, defining $\bar{A}(\mu) = \theta_1(\mu)\bar{A}_1 + \dots + \theta_Q(\mu)\bar{A}_Q$, such that*

$$\lambda_{\min}(\bar{A}(\tilde{\mu})) = \lambda_{\text{LB}}(\tilde{\mu}; C_J) \quad \text{and} \quad \lambda_{\min}(\bar{A}(\mu_i)) = \lambda_{\min}(A(\mu_i)) \quad (10)$$

hold for $i = 1, \dots, J$.

Proof. Let the columns of $V \in \mathbb{C}^{N \times J}$ and $V_{\perp} \in \mathbb{C}^{N \times (N-J)}$ form orthonormal bases of $\mathcal{V} = \text{span}\{v_1, \dots, v_J\}$ and \mathcal{V}^{\perp} , respectively, where each v_i denotes an eigenvector associated with $\lambda_{\min}(A(\mu_i))$. Moreover, let $y_{\tilde{\mu}} \in \mathcal{Y}_{\text{LB}}(C_J) \subset \mathbb{R}^Q$ denote a minimizer of (7) for $\tilde{\mu}$, that is, $\lambda_{\text{LB}}(\tilde{\mu}; C_J) = \theta(\tilde{\mu})^T y_{\tilde{\mu}}$. The rest of the proof consists of showing that the matrices defined by

$$\bar{A}_q := VV^*A_qVV^* + y_{\tilde{\mu},q}V_{\perp}V_{\perp}^*, \quad q \in \{1, \dots, Q\},$$

satisfy (10).

Given $u \in \mathbb{C}^N$, we can write $u = u_{\mathcal{V}} + u_{\perp}$ with $u_{\mathcal{V}} \in \mathcal{V}$ and $u_{\perp} \in \mathcal{V}^{\perp}$. For any $\mu \in D$, we therefore have

$$u^* \bar{A}(\mu) u = \sum_{q=1}^Q \theta_q(\mu) (u_{\mathcal{V}}^* A_q u_{\mathcal{V}} + y_{\tilde{\mu}, q} \|u_{\perp}\|_2^2) = u_{\mathcal{V}}^* A(\mu) u_{\mathcal{V}} + \theta(\mu)^T y_{\tilde{\mu}} \|u_{\perp}\|_2^2. \quad (11)$$

For $\mu = \mu_i$, this yields

$$u^* \bar{A}(\mu_i) u \geq \lambda_{\min}(A(\mu_i)) \|u_{\mathcal{V}}\|_2^2 + \lambda_{\min}(A(\mu_i)) \|u_{\perp}\|_2^2 = \lambda_{\min}(A(\mu_i)) \|u\|_2^2,$$

where we used that $y_{\tilde{\mu}} \in \mathcal{Y}_{\text{LB}}(C_J)$ implies $\theta(\mu_i)^T y_{\tilde{\mu}} \geq \lambda_{\min}(A(\mu_i))$. Since equality is attained for $u = v_i$, this implies the second equality in (10).

To show the first equality, we first notice that the definition of $\lambda_{\text{LB}}(\tilde{\mu}; C_J)$ implies

$$u_{\mathcal{V}}^* A(\tilde{\mu}) u_{\mathcal{V}} = \theta(\tilde{\mu})^T R(u_{\mathcal{V}}) \|u_{\mathcal{V}}\|_2^2 \geq \lambda_{\text{LB}}(\tilde{\mu}; C_J) \|u_{\mathcal{V}}\|_2^2.$$

Inserted into (11) for $\mu = \tilde{\mu}$, this yields

$$u^* \bar{A}(\tilde{\mu}) u \geq \lambda_{\text{LB}}(\tilde{\mu}; C_J) \|u_{\mathcal{V}}\|_2^2 + \lambda_{\text{LB}}(\tilde{\mu}; C_J) \|u_{\perp}\|_2^2 = \lambda_{\text{LB}}(\tilde{\mu}; C_J) \|u\|_2^2.$$

Since equality is attained by any $u \in \mathcal{V}^{\perp}$, this shows the first equality in (10) and thus completes the proof. \square

Since the definition of the lower bounds in (7) only depends on $\theta(\mu)$ and the eigenvalues at μ_i , the lower bounds for the matrix function $\bar{A}(\mu)$ constructed in Theorem 2.3 are identical with those for $A(\mu)$. For $\bar{A}(\mu)$, the lower bound $\lambda_{\text{LB}}(\tilde{\mu}; C_J)$ coincides with the exact eigenvalue at an arbitrary fixed $\tilde{\mu} \in D$. Hence, additional knowledge, beyond the eigenvalues at μ_i , needs to be incorporated to improve the lower bounds.

3 Subspace approach

In this section, our new subspace approach is presented that takes eigenvector information across different parameter samples into account and offers the flexibility to incorporate eigenvectors for larger eigenvalues as well.

Given $C_J = \{\mu_1, \dots, \mu_J\} \subset D$, suppose that for each sample μ_i we have computed the $\ell \geq 1$ smallest eigenvalues

$$\lambda_i = \lambda_i^{(1)} \leq \lambda_i^{(2)} \leq \dots \leq \lambda_i^{(\ell)}$$

of $A(\mu_i)$ along with an orthonormal basis of associated eigenvectors $v_i^{(1)}, v_i^{(2)}, \dots, v_i^{(\ell)} \in \mathbb{C}^N$. To simplify notation, we assume that an equal number of eigenpairs has been computed for each μ_1, \dots, μ_J , although this is not necessary. The eigenvectors will be collected in the subspace

$$\mathcal{V}(C_J, \ell) := \text{span}\{v_1^{(1)}, \dots, v_1^{(\ell)}, v_2^{(1)}, \dots, v_2^{(\ell)}, \dots, v_J^{(1)}, \dots, v_J^{(\ell)}\}. \quad (12)$$

In the subsequent two sections, we discuss how the information in $\mathcal{V}(C_J, \ell)$ can be used to compute tighter bounds for $\lambda_{\min}(A(\mu))$.

3.1 Subspace approach for upper bounds

Given the subspace $\mathcal{V}(C_J, \ell)$ from (12), we define an upper bound set analogously to (5):

$$\mathcal{Y}_{\text{SUB}}(C_J, \ell) := \{R(v) : v \in \mathcal{V}(C_J, \ell)\}.$$

The corresponding upper bound for $\mu \in D$ is defined as

$$\lambda_{\text{SUB}}(\mu; C_J, \ell) := \min_{y \in \mathcal{Y}_{\text{SUB}}(C_J, \ell)} \theta(\mu)^T y.$$

Clearly, we have $\mathcal{Y}_{\text{UB}}(C_J) \subseteq \mathcal{Y}_{\text{SUB}}(C_J, \ell) \subseteq \mathcal{Y}$ and thus

$$\lambda_{\text{UB}}(\mu; C_J) \geq \lambda_{\text{SUB}}(\mu; C_J, \ell) \geq \lambda_{\min}(A(\mu)).$$

To evaluate $\lambda_{\text{SUB}}(\mu; C_J, \ell)$, we first compute an orthonormal basis $V \in \mathbb{C}^{N \times J\ell}$ of $\mathcal{V}(C_J, \ell)$ and obtain

$$\begin{aligned} \lambda_{\text{SUB}}(\mu; C_J, \ell) &= \min_{v \in \mathcal{V}(C_J, \ell)} \theta(\mu)^T R(v) = \min_{\substack{w \in \mathbb{C}^{J\ell} \\ \|w\|_2=1}} \theta(\mu)^T R(Vw) \\ &= \min_{\substack{w \in \mathbb{C}^{J\ell} \\ \|w\|_2=1}} \theta_1(\mu) w^* V^* A_1 V w + \dots + \theta_Q(\mu) w^* V^* A_Q V w \\ &= \lambda_{\min}(\theta_1(\mu) V^* A_1 V + \dots + \theta_Q(\mu) V^* A_Q V) = \lambda_{\min}(V^* A(\mu) V). \end{aligned} \quad (13)$$

Thus, the computation of $\lambda_{\text{SUB}}(\mu, C_J, \ell)$ requires the solution of an eigenvalue problem of size $J\ell \times J\ell$, with $J\ell$ usually much smaller than N .

3.2 Subspace approach for lower bounds

We will use a perturbation result to turn the upper bound (13) into a lower bound $\lambda_{\text{SLB}}(\mu; C_J, \ell)$ for $\mu \in D$. For this purpose, we consider for some small integer $r \leq J\ell$ the r smallest eigenvalues

$$\lambda_{\text{SUB}}(\mu; C_J, \ell) = \lambda_{\mathcal{V}}^{(1)} \leq \lambda_{\mathcal{V}}^{(2)} \leq \dots \leq \lambda_{\mathcal{V}}^{(r)}$$

of $V^* A(\mu) V$, along with the corresponding eigenvectors $w_1, \dots, w_r \in \mathbb{C}^{J\ell}$. Let $U \in \mathbb{C}^{N \times r}$ be an orthonormal basis of the subspace $\mathcal{U}(\mu)$ spanned by the Ritz vectors:

$$\mathcal{U}(\mu) := \text{span}\{Vw_1, \dots, Vw_r\}.$$

Moreover, let $U_{\perp} \in \mathbb{C}^{N \times (N-r)}$ be an orthonormal basis of $\mathcal{U}^{\perp}(\mu)$ and denote the eigenvalues of $U_{\perp}^* A(\mu) U_{\perp}$ by

$$\lambda_{\mathcal{U}^{\perp}}^{(1)} \leq \lambda_{\mathcal{U}^{\perp}}^{(2)} \leq \dots \leq \lambda_{\mathcal{U}^{\perp}}^{(N-r)}.$$

The transformed matrix

$$[U, U_{\perp}]^* A(\mu) [U, U_{\perp}] = \begin{bmatrix} U^* A(\mu) U & U^* A(\mu) U_{\perp} \\ U_{\perp}^* A(\mu) U & U_{\perp}^* A(\mu) U_{\perp} \end{bmatrix}$$

clearly has the same eigenvalues as $A(\mu)$, while the perturbed matrix

$$\begin{bmatrix} U^* A(\mu) U & 0 \\ 0 & U_{\perp}^* A(\mu) U_{\perp} \end{bmatrix}$$

has the eigenvalues $\{\lambda_{\mathcal{V}}^{(1)}, \dots, \lambda_{\mathcal{V}}^{(r)}\} \cup \{\lambda_{\mathcal{U}^\perp}^{(1)}, \dots, \lambda_{\mathcal{U}^\perp}^{(N-r)}\}$. Applying a perturbation result by Li and Li [22] to this situation yields the error bound

$$|\lambda_{\min}(A(\mu)) - \min(\lambda_{\mathcal{V}}^{(1)}, \lambda_{\mathcal{U}^\perp}^{(1)})| \leq \frac{2\rho^2}{\delta + \sqrt{\delta^2 + 4\rho^2}},$$

with the residual norm

$$\rho := \|U_\perp^* A(\mu) U\|_2 = \|A(\mu) U - U(U^* A(\mu) U)\|_2$$

and the absolute gap $\delta := |\lambda_{\mathcal{V}}^{(1)} - \lambda_{\mathcal{U}^\perp}^{(1)}|$. Rearranging terms thus gives the lower bound

$$f(\lambda_{\mathcal{U}^\perp}^{(1)}) \leq \lambda_{\min}(A(\mu)), \quad \text{with} \quad f(\eta) := \min(\lambda_{\mathcal{V}}^{(1)}, \eta) - \frac{2\rho^2}{|\lambda_{\mathcal{V}}^{(1)} - \eta| + \sqrt{|\lambda_{\mathcal{V}}^{(1)} - \eta|^2 + 4\rho^2}}. \quad (14)$$

This lower bound is not practical so far, as it involves the quantity $\lambda_{\mathcal{U}^\perp}^{(1)}$, which would require the solution of a large eigenvalue problem of size $(N-r) \times (N-r)$.

Lemma 3.1. *The function $f : \mathbb{R} \rightarrow \mathbb{R}$ defined in (14) is continuous and monotonically increasing.*

Proof. See Section A. □

Lemma 3.1 implies that $f(\eta)$ remains a lower bound as long as $\eta \leq \lambda_{\mathcal{U}^\perp}^{(1)}$. To summarize, our subspace-accelerated lower bound is defined as

$$\lambda_{\text{SLB}}(\mu; C_J, \ell) := \min(\lambda_{\mathcal{V}}^{(1)}, \eta) - \frac{2\rho^2}{|\lambda_{\mathcal{V}}^{(1)} - \eta| + \sqrt{|\lambda_{\mathcal{V}}^{(1)} - \eta|^2 + 4\rho^2}} \quad (15)$$

for a lower bound η of $\lambda_{\mathcal{U}^\perp}^{(1)}$.

3.2.1 Determining a lower bound for $\lambda_{\mathcal{U}^\perp}^{(1)}$

The lower bound for $\lambda_{\mathcal{U}^\perp}^{(1)} = \lambda_{\min}(U_\perp^* A(\mu) U_\perp)$ needed in (15) will be determined by adapting the ideas from Section 2.1. Let us recall that SCM determines a lower bound for $\lambda_{\min}(A(\mu))$ by solving the LP

$$\lambda_{\text{LB}}(\mu; C_J) = \min_{y \in \mathcal{Y}_{\text{LB}}(C_J)} \theta(\mu)^T y, \quad (16)$$

with $\mathcal{Y}_{\text{LB}}(C_J) := \{y \in \mathcal{B} : \theta(\mu_i)^T y \geq \lambda_i, i = 1, \dots, J\}$ and the bounding box \mathcal{B} defined in (6). To simplify the discussion, we always assume in the following that $\mathcal{Y}_{\text{LB}}(C_J)$ is a simple polytope with no degenerate facets. Then there exists an optimizer $y_\mu \in \mathbb{R}^Q$ of (16) such that there are Q , among $2Q+J$, linearly independent active constraints [26]. In other words, y_μ satisfies a linear system

$$\Theta y_\mu = \psi, \quad (17)$$

where $\Theta \in \mathbb{R}^{Q \times Q}$ is invertible and each equation corresponds either to a constraint of the form $\theta(\mu_i)^T y_\mu = \lambda_i$ or to a box constraint. In the following, we tacitly assume that at least one of the active constraints is a non-box constraint.

Establishing a lower bound for $\lambda_{\mathcal{U}_\perp}^{(1)}$ is equivalent to determining η such that $\eta \leq u_\perp^* A(\mu) u_\perp$ holds for every $u_\perp \in \mathcal{U}_\perp(\mu)$ with $\|u_\perp\|_2 = 1$. The restriction of u_\perp to a lower-dimensional subspace can be used to tighten the non-box constraints in (16).

Lemma 3.2. *With the notation introduced above, let $\Lambda_i = \text{diag}(\lambda_i^{(1)}, \dots, \lambda_i^{(\ell)})$ and $V_i = [v_i^{(1)}, \dots, v_i^{(\ell)}]$. If $N - r \geq r$ then*

$$u_\perp^* A(\mu_i) u_\perp \geq \lambda_i + \beta_i,$$

where β_i is the smallest eigenvalue of the matrix

$$(\Lambda_i - \lambda_i I_\ell) - V_i^* U U^* V_i (\Lambda_i - \lambda_i^{(\ell+1)} I_\ell).$$

Proof. Using the spectral decomposition of $A(\mu_i)$, the result follows from

$$\begin{aligned} \min_{\substack{u_\perp \in \mathcal{U}_\perp \\ \|u_\perp\|_2=1}} u_\perp^* A(\mu_i) u_\perp &\geq \min_{\substack{u_\perp \in \mathcal{U}_\perp \\ \|u_\perp\|_2=1}} u_\perp^* V_i \Lambda_i V_i^* u_\perp + \lambda_i^{(\ell+1)} u_\perp^* (I - V_i V_i^*) u_\perp & (18) \\ &= \lambda_i^{(\ell+1)} + \min_{\substack{u_\perp \in \mathcal{U}_\perp \\ \|u_\perp\|_2=1}} u_\perp^* V_i (\Lambda_i - \lambda_i^{(\ell+1)} I_\ell) V_i^* u_\perp \\ &= \lambda_i^{(\ell+1)} + \lambda_{\min}(U_\perp^* V_i (\Lambda_i - \lambda_i^{(\ell+1)} I_\ell) V_i^* U_\perp) \\ &= \lambda_i^{(\ell+1)} + \lambda_{\min}(V_i^* U_\perp U_\perp^* V_i (\Lambda_i - \lambda_i^{(\ell+1)} I_\ell)) \\ &= \lambda_i^{(\ell+1)} + \lambda_{\min}((I_\ell - V_i^* U U^* V_i) (\Lambda_i - \lambda_i^{(\ell+1)} I_\ell)) \\ &= \lambda_i + \lambda_{\min}((\Lambda_i - \lambda_i I_\ell) - V_i^* U U^* V_i (\Lambda_i - \lambda_i^{(\ell+1)} I_\ell)), \end{aligned}$$

where we used in the third equality that the negative eigenvalues of the matrix product $U_\perp^* V_i (\Lambda_i - \lambda_i^{(\ell+1)} I_\ell) V_i^* U_\perp$ do not change under a cyclic permutation of its factors. \square

Using the values of β_i defined in Lemma 3.2, we update the right-hand side $\psi \in \mathbb{R}^Q$ in (17) as follows: If the k th equation corresponds to a non-box constraint $\theta(\mu_i)^T y = \lambda_i$, we set $\tilde{\psi}_k := \psi_k + \beta_i = \lambda_i + \beta_i$ and, otherwise, $\tilde{\psi}_k := \psi_k$. Since Θ is invertible, the solution of the resulting LP

$$\inf_y \theta(\mu)^T y \quad \text{subject to} \quad \Theta y \geq \tilde{\psi}$$

is trivially given by

$$\tilde{y}_\mu := \Theta^{-1} \tilde{\psi}. \quad (19)$$

This finally yields the desired lower bound

$$\eta(\mu) := \theta(\mu)^T \tilde{y}_\mu \leq \lambda_{\mathcal{U}_\perp}^{(1)} = \lambda_{\min}(U_\perp^* A(\mu) U_\perp).$$

Remark 3.3. *The choice of r , the dimension of the Ritz subspace $\mathcal{U}(\mu)$, requires some consideration. For $r = 0$, $\mathcal{U}_\perp(\mu) = \mathbb{R}^N$ yields no improvement: $\lambda_{\text{SLB}}(\mu; C_J, \ell) = \lambda_{\text{LB}}(\mu; C_J, \ell)$. Intuitively, choosing $r = 1$ will be most effective when the second smallest eigenvalue of $A(\mu)$ is well separated from the smallest eigenvalue. Otherwise, one may benefit from choosing slightly larger values of r . In practice, we choose r by taking the maximal value of $\lambda_{\text{SLB}}(\mu; C_J, \ell)$ for a few small values of $r = 0, 1, 2, \dots$.*

3.3 Algorithm and computational complexity

The implementation of the proposed lower and upper bounds requires some care in order to avoid unnecessary that involve quantities of size N , the original size of the problem.

Computation of ρ . The quantity $\rho = \|A(\mu)U - U\Lambda_U\|_2$ with $\Lambda_U = U^*A(\mu)U = \text{diag}(\lambda_{\mathcal{V}}^{(1)}, \dots, \lambda_{\mathcal{V}}^{(r)})$ can be computed by solving an $r \times r$ eigenvalue problem:

$$\begin{aligned} \rho^2 &= \lambda_{\max}((A(\mu)U(\mu) - U(\mu)\Lambda(\mu))^*(A(\mu)U(\mu) - U(\mu)\Lambda(\mu))) \\ &= \lambda_{\max}(U(\mu)^*A(\mu)^*A(\mu)U(\mu) - \Lambda(\mu)^2). \end{aligned}$$

Computation of $V^*A(\mu)V$ and $U^*A(\mu)^*A(\mu)U$. By the affine linear decomposition (2),

$$V^*A(\mu)V = \theta_1(\mu)V^*A_1V + \dots + \theta_Q(\mu)V^*A_QV.$$

A standard technique in RBM, we compute and store the $J\ell \times J\ell$ matrices V^*A_qV , and update them as new columns are added to V . In turn, the computation of $V^*A(\mu)V$, which is needed to evaluate the upper bound for every $\mu \in \Xi$, becomes negligible as long as $J\ell \ll N$. Similarly, the evaluation of $U^*A(\mu)^*A(\mu)U$ needed for ρ becomes negligible after the precomputation of $V^*A_q^*A_{q'}V$ for all $q, q' = 1, \dots, Q$.

Choice of ℓ . Clearly, a larger choice of ℓ can be expected to lead to better bounds. On the other hand, a larger value of ℓ increases the computational cost. Intuitively, choosing ℓ larger than one appears to be most beneficial when the gap between the smallest and second smallest eigenvalues is small or even vanishes. One could, for example, choose ℓ such that $\lambda_i^{(\ell+1)} - \lambda_i^{(1)}$ exceeds a certain threshold. However, in the absence of a priori information on eigenvalue gaps, it might be wisest to simply choose $\ell = 1$ for all μ_i .

Algorithm 2 summarizes our proposed procedure for computing subspace lower and upper bounds. Similarly as SCM, the algorithm requires the solution of $2Q + J$ eigenvalue problems of size $N \times N$ for determining the bounding box in the beginning and the smallest $\ell + 1$ eigenpairs in each iteration. Clearly, the latter part will become more expensive if $\ell \geq 1$. However, we expect that this increase can be mitigated significantly in practice by the use of block algorithms. More specifically, when using a block eigenvalue solver such as LOBPCG [17] and efficient implementations of block matrix-vector products with the matrix A (and its preconditioner), the computation of ℓ smallest eigenvalues will not be much more expensive as long as ℓ remains modest.

Computing $\lambda_{\text{SUB}}(\mu; C_J, \ell)$ and $\lambda_{\text{SLB}}(\mu; C_J, \ell)$ for all $\mu \in \Xi$ amounts to solving $J|\Xi|$ LP problems with Q variables and $2Q + J$ constraints, as well as $J|\Xi|$ eigenproblems of size $J\ell \times J\ell$. As long as $J\ell \ll N$, these parts will be negligible, and the cost of Algorithms 1 and 2 will be approximately equal.

3.4 Interpolation results

By definition, we already know that the bounds from our subspace approach are never worse than the bounds produced by SCM:

$$\lambda_{\text{LB}}(\mu; C_J) \leq \lambda_{\text{SLB}}(\mu; C_J, \ell) \leq \lambda_{\min}(A(\mu)) \leq \lambda_{\text{SUB}}(\mu; C_J, \ell) \leq \lambda_{\text{UB}}(\mu; C_J), \quad (20)$$

Algorithm 2 Subspace SCM

Input: Training set Ξ , affine linear decomposition such that $A(\mu) = \theta_1(\mu)A_1 + \dots + \theta_Q(\mu)A_Q$ is Hermitian for every $\mu \in \Xi$. Relative error tolerance ε_{SCM} .

Output: Set $C_J \subset \Xi$ with corresponding eigenvalues $\lambda_i^{(j)}$ and eigenvector basis V of $\mathcal{V}(C_J, \ell)$, such that $\frac{\lambda_{\text{SUB}}(\mu; C_J, \ell) - \lambda_{\text{SLB}}(\mu; C_J, \ell)}{\lambda_{\text{SUB}}(\mu; C_J, \ell)} < \varepsilon_{\text{SCM}}$ for every $\mu \in \Xi$.

- 1: compute $\lambda_{\min}(A_q), \lambda_{\max}(A_q)$ for $q = 1, \dots, Q$, defining \mathcal{B} according to (6)
 - 2: $J = 0, C_0 = \emptyset$
 - 3: **while** $\max_{\mu \in \Xi} \frac{\lambda_{\text{SUB}}(\mu; C_J, \ell) - \lambda_{\text{SLB}}(\mu; C_J, \ell)}{\lambda_{\text{SUB}}(\mu; C_J, \ell)} > \varepsilon_{\text{SCM}}$ **do**
 - 4: $\mu_{J+1} \leftarrow \arg \max_{\mu \in \Xi} \frac{\lambda_{\text{SUB}}(\mu; C_J, \ell) - \lambda_{\text{SLB}}(\mu; C_J, \ell)}{\lambda_{\text{SUB}}(\mu; C_J, \ell)}$
 - 5: compute smallest eigenpairs $(\lambda_{J+1}^{(1)}, v_{J+1}^{(1)}), \dots, (\lambda_{J+1}^{(\ell)}, v_{J+1}^{(\ell)})$ of $A(\mu_{J+1})$
 - 6: $C_{J+1} \leftarrow C_J \cup \mu_{J+1}$
 - 7: update $V^* A_q V$ and $V^* A_q^* A_{q'} V$ for all $q, q' = 1, \dots, Q$
 - 8: **for** $\mu \in \Xi$ **do**
 - 9: compute $\lambda_{\text{SUB}}(\mu; C_{J+1}, \ell) = \lambda_{\min}(V^* A(\mu) V)$
 - 10: compute $\rho = \sqrt{\lambda_{\max}(U(\mu)^* A(\mu)^* A(\mu) U(\mu) - \Lambda(\mu)^2)}$
 - 11: compute $y_\mu = \arg \min_{y \in \mathcal{Y}_{\text{LB}}(C_{J+1})} \theta(\mu)^T y$ and updated \tilde{y}_μ according to (19)
 - 12: $\eta(\mu) \leftarrow \theta(\mu)^T \tilde{y}_\mu$
 - 13: compute $\lambda_{\text{SLB}}(\mu; C_{J+1}, \ell)$ according to (15)
 - 14: **end for**
 - 15: $J \leftarrow J + 1$
 - 16: **end while**
-

with equality at $\mu = \mu_i \in C_J$. Together with Theorem 2.1, these inequalities imply that our upper bounds also interpolate the derivatives at μ_i .

Corollary 3.4. *For any $\ell \geq 1$ and any $\mu_i \in C_J$ that satisfies the assumptions of Theorem 2.1, it holds that*

$$\nabla \lambda_{\text{SUB}}(\mu_i; C_J, \ell) = \nabla \lambda_{\min}(A(\mu_i))$$

with $\lambda_{\text{SUB}}(\mu_i; C_J, \ell)$ defined as in (13).

Proof. By the assumptions, μ_i is an interior point of D and (20) implies that the inequality $\lambda_{\min}(A(\mu)) \leq \lambda_{\text{SUB}}(\mu; C_J, \ell) \leq \lambda_{\text{UB}}(\mu; C_J)$ holds for all μ in a neighbourhood of μ_i . Combined with the result $\nabla \lambda_{\min}(A(\mu_i)) = \nabla \lambda_{\text{UB}}(\mu_i; C_J)$ of Theorem 2.1, this implies $\nabla \lambda_{\text{SUB}}(\mu_i; C_J, \ell) = \nabla \lambda_{\min}(A(\mu_i))$. \square

In contrast to SCM, it turns out that the subspace lower bounds also interpolate the derivative of $\lambda_{\min}(A(\mu))$ at $\mu \in C_J$. To show this, we need the following lemma.

Lemma 3.5. *Let $\mu_i \in C_J$ satisfy the assumptions of Theorem 2.1. For any $\varepsilon > 0$, there is a neighbourhood $\Omega \subseteq D$ around μ_i such that*

$$|\lambda_i - \lambda_{\mathcal{V}}^{(1)}(\mu)| \leq \varepsilon, \tag{21}$$

$$\lambda_i^{(2)} - \eta(\mu) \leq \varepsilon, \tag{22}$$

hold for all $\mu \in \Omega$, where $\lambda_{\mathcal{V}}^{(1)}(\cdot)$ and $\eta(\cdot)$ are defined as in Section 3.2.

Proof. By construction, $\lambda_{\mathcal{V}}^{(1)}(\mu_i) = \lambda_i$ and thus the continuity of the smallest eigenvalue implies that (21) holds for all μ in some neighbourhood Ω_1 around μ_i . It remains to prove (22).

In the LP (7) for determining $\lambda_{\text{LB}}(\mu_i; C_J)$, which is trivially given by λ_i , the constraint $\theta(\mu_i)^T y = \lambda_i$ is active. Since we assumed that $\mathcal{Y}_{\text{LB}}(C_J)$ is a simple polytope with no degenerate facets, the continuity of $\theta(\mu)$ implies that this constraint remains active in a neighbourhood Ω_2 : $\theta(\mu_i)^T y_\mu = \lambda_i$ for all $\mu \in \Omega_2$, where y_μ is a minimizer of (7) for determining $\lambda_{\text{LB}}(\mu; C_J)$.

By (18), the value of β_i defined in Lemma 3.2 satisfies

$$\beta_i = \min_{\substack{u_\perp \in \mathcal{U}^\perp(\mu) \\ \|u_\perp\|_2=1}} u_\perp^* V_i \Lambda_i V_i^* u_\perp + \lambda_i^{(\ell+1)} u_\perp^* (I - V_i V_i^*) u_\perp - \lambda_i.$$

The eigenvector $v_i^{(1)}$ belonging to the eigenvalue $\lambda_i = \lambda_{\min}(A(\mu_i))$ is contained in \mathcal{U} for $\mu = \mu_i$. Once again the simplicity of λ_i implies that the angle between $v_i^{(1)}$ and \mathcal{U} becomes arbitrarily small as μ approaches μ_i . Therefore, for any $\varepsilon > 0$, there is a neighbourhood Ω_3 such that

$$\beta_i \geq \min_{\substack{u_\perp \perp v_i^{(1)} \\ \|u_\perp\|_2=1}} u_\perp^* V_i \Lambda_i V_i^* u_\perp + \lambda_i^{(\ell+1)} u_\perp^* (I - V_i V_i^*) u_\perp - \lambda_i - \frac{\varepsilon}{2} = \lambda_i^{(2)} - \lambda_i - \frac{\varepsilon}{2}.$$

In summary, the vector \check{y}_μ defined in (19) satisfies

$$\theta(\mu_i)^T \check{y}_\mu = \lambda_i + \beta_i \geq \lambda_i + \lambda_i^{(2)} - \lambda_i - \frac{\varepsilon}{2} = \lambda_i^{(2)} - \frac{\varepsilon}{2}. \quad (23)$$

By the invertibility of Θ , the vector \check{y}_μ remains bounded in the vicinity of μ_i . Together with the continuity of $\theta(\mu)$, this implies that there is a neighbourhood Ω_4 such that

$$|(\theta(\mu) - \theta(\mu_i))^T \check{y}_\mu| \leq \frac{\varepsilon}{2}, \quad \forall \mu \in \Omega_4.$$

Combined with (23), this yields

$$\eta(\mu) = \theta(\mu)^T \check{y}_\mu \geq \lambda_i^{(2)} - \varepsilon,$$

which establishes (22). Setting $\Omega = \Omega_1 \cap \Omega_2 \cap \Omega_3 \cap \Omega_4$ completes the proof. \square

The following theorem establishes the Hermite interpolation property of the subspace lower bounds.

Theorem 3.6. *Let $\mu_i \in C_J$ satisfy the assumptions of Theorem 2.1 and, additionally, suppose that $r \leq \ell$ and $\lambda_i^{(r+1)} > \lambda_i^{(r)}$. Then*

$$\nabla \lambda_{\text{SLB}}(\mu_i; C_J, \ell) = \nabla \lambda_{\min}(A(\mu_i)).$$

Proof. By Lemma 3.5 and the simplicity of $\lambda_{\min}(A(\mu_i))$, there is $\delta_0 > 0$ such that $\eta(\mu) \geq \lambda_{\mathcal{V}}^{(1)}(\mu) + \delta_0$ for μ sufficiently close to μ_i . Hence, the subspace lower bound (15) is given by

$$\lambda_{\text{SLB}}(\mu; C_J, \ell) = \lambda_{\mathcal{V}}^{(1)}(\mu) - \frac{2\rho^2}{\delta + \sqrt{\delta^2 + 4\rho^2}}, \quad (24)$$

with $\rho = \|U^\perp A(\mu)U\|_2$ and $\delta = |\lambda_{\mathcal{V}}^{(1)}(\mu) - \eta(\mu)| \geq \delta_0$. By Corollary 3.4, we have

$$\nabla \lambda_{\mathcal{V}}^{(1)}(\mu_i) = \nabla \lambda_{\text{SUB}}(\mu_i; C_J, \ell) = \nabla \lambda_{\min}(A(\mu_i)).$$

Since δ is bounded from below, the result follows from (24) if the gradient of ρ^2 is zero. The assumptions $\lambda_i^{(r+1)} > \lambda_i^{(r)}$ and $r \leq \ell$ imply that the invariant subspace belonging to the r smallest eigenvalues of $A(\mu_i)$ is simple and contained in V . Hence, standard perturbation results for invariant subspaces [36] yield $\rho = O(\|\mu - \mu_i\|_2)$ for $\mu \rightarrow \mu_i$ and therefore $\nabla \rho^2 = 0$, which completes the proof. \square

3.4.1 A priori convergence in the one-parameter case

In the following, we analyse the convergence of our subspace bounds for a special case: We assume that $A(\mu)$ depends analytically on one parameter $\mu \in [-1, 1]$ and, moreover, the eigenvalue $\lambda_{\min}(A(\mu))$ is simple for all $\mu \in [-1, 1]$.

Let E_R denote the open elliptic disc with foci ± 1 and the sum of its half axes equal to R . Under the above assumptions, there is $R_0 > 1$ such that the (suitably normalized) eigenvector $v(\mu)$ belonging to $\lambda_{\min}(A(\mu))$ admits an analytic extension $v : E_{R_0} \rightarrow \mathbb{C}^N$; see, e.g., [15, 32]. Note that v can be chosen to have norm 1 on $[-1, 1]$, see [32, Theorem XII.4], but this is not the case on E_{R_0} in general. Let $C_J = \{\mu_1, \dots, \mu_J\}$ contain the Chebyshev nodes $\mu_i = \cos(\frac{2i-1}{2J}\pi)$ and set $v_i := v(\mu_i)$. The corresponding vector-valued interpolating polynomial is given by

$$p_J(\mu) = \ell_1(\mu)v_1 + \dots + \ell_J(\mu)v_J \quad (25)$$

with the Lagrange polynomials $\ell_1, \dots, \ell_J : [-1, 1] \rightarrow \mathbb{R}$. By extending a standard interpolation error result [25, Corollary 6.6A] to vector-valued functions (see, e.g., [18, Lemma 2.2] for a similar extension), one can show that

$$\max_{\mu \in [-1, 1]} \|v(\mu) - p_J(\mu)\|_2 \leq \frac{(R + R^{-1})M}{(R^{J+1} - R^{-J-1})(R + R^{-1} - 2)} \quad (26)$$

holds for any $1 < R < R_0$, with $M = \sup_{z \in E_R} \|v(z)\|_2$. This result is utilized in the proof of the following theorem, which shows exponential convergence of our subspace bounds.

Theorem 3.7. *Under the setting described above, the subspace lower and upper bounds for $\ell = r = 1$ satisfy*

$$\lambda_{\text{SUB}}(\mu; C_J, 1) - \lambda_{\min}(A(\mu)) \leq C_U R^{-2J} \quad (27)$$

$$\lambda_{\min}(A(\mu)) - \lambda_{\text{SLB}}(\mu; C_J, 1) \leq C_L R^{-2J} \quad (28)$$

for every $\mu \in [-1, 1]$, with constants C_U, C_L independent of J and μ .

Proof. For $\ell = 1$, the subspace used in our bounds takes the form $\mathcal{V} = \text{span}\{v_1, \dots, v_J\}$. The interpolating polynomial defined in (25) clearly satisfies $p_J(\mu) \in \mathcal{V}$ and hence (26) yields the following bound on the angle between \mathcal{V} and $v(\mu)$:

$$\min_{\tilde{v} \in \mathcal{V}} \|\tilde{v} - v(\mu)\|_2 \lesssim R^{-J}. \quad (29)$$

The inequality (27) now follows from existing approximation results for Ritz values; see, e.g., [29, Chapter 11].

To prove (28), we first note that the arguments from the proof of Theorem 3.6 can be utilized to show that

$$\lambda_{\text{SLB}}(\mu; C_J, 1) = \lambda_{\text{SUB}}(\mu; C_J, 1) - \frac{2\rho^2}{\delta + \sqrt{\delta^2 + 4\rho^2}},$$

for sufficiently large J , where $\delta > \delta_0 > 0$ for some δ_0 not depending on μ or J . Since $r = 1$, the quantity ρ coincides with the residual of the smallest Ritz vector of $A(\mu)$ with respect to \mathcal{V} . By (29) and [29, Theorem 11.7.1], we have $\rho \lesssim R^{-J}$, which completes the proof. \square

The maximal value of the exponent R in (27)–(28) depends on the gap between the smallest and the second smallest eigenvalue and the variation of $A(\mu)$. This is discussed in more detail for a special case in [1, Section 2.3.1].

3.5 Residual-based lower bounds

As we will see in the numerical experiments (especially in Example 4.5), the subspace lower bounds can sometimes converge rather slowly in the initial phase of the algorithm, in contrast to the subspace upper bounds. This slow convergence can be viewed as a price that needs to be paid in order to maintain the reliability of the lower bounds. In the following, we will propose an alternative that is heuristic (i.e., its reliability is not guaranteed) and is observed to converge faster in the initial phase.

The alternative consists of simply subtracting the residual norm from the upper bound:

$$\lambda_{\text{SUB}}(\mu; C_J, \ell) - \|A(\mu)u - \lambda_{\text{SUB}}(\mu; C_J, \ell)u\|_2, \quad (30)$$

where u with $\|u\|_2 = 1$ is a Ritz vector belonging to the smallest Ritz value $\lambda_{\text{SUB}}(\mu; C_J, \ell)$ of $A(\mu)$ with respect to \mathcal{V} . A basic first-order perturbation result for Hermitian matrices [29] implies that (30) constitutes a lower bound for *an* eigenvalue of $A(\mu)$, but not necessarily the smallest one. There is a risk, especially in the very beginning, that (30) is actually larger than the smallest eigenvalue, see Section 4 for examples. However, in all numerical experiments we have observed that a small number of iterations suffices until (30) becomes a lower bound for the smallest eigenvalue.

Remark 3.8. *When using the residual-based lower bound (30), it makes sense to also adjust the error measure (8) that drives the sampling strategy to*

$$\max_{\mu \in \Xi} \frac{\|A(\mu)u - \lambda_{\text{SUB}}(\mu; C_J, \ell)u\|_2}{|\lambda_{\text{SUB}}(\mu; C_J, \ell)|},$$

and stop the iteration when this error estimate drops below ε_{SCM} .

4 Applications and numerical experiments

In this section, we report on the performance of our proposed approach for a number of large-scale examples. Algorithms 1 and 2 have been implemented in MATLAB Version

7.14.0.739 (R2012a) and all experiments have been performed on an Intel Xeon CPU E31225 with 4 cores, 3.1 GHz, and 8 GB RAM.

We compare Algorithm 2 with Algorithm 1 by computing the maximum relative error ratio (8). Additionally, we compare the convergence of the bounds from Sections 2 and 3 towards the exact smallest eigenvalues by measuring the absolute error

$$\max_{\mu \in \Xi} |\text{bound}(\mu) - \lambda_{\min}(A(\mu))|, \quad (31)$$

for the corresponding bound, both with respect to the number of iterations and with respect to the execution time (in seconds).

When implementing and testing Algorithms 1 and 2, we have made the following choices. We set the relative tolerance to $\varepsilon_{SCM} = 10^{-4}$, the maximum number of iterations to $J_{\max} = 200$ and the surrogate set Ξ to be a random subset of D containing 1000 elements. The smallest eigenpairs of $A(\mu_i)$ have been computed using the MATLAB built-in function `eigs`, which is based on ARPACK [20], with the tolerance set to 10^{-6} . The MATLAB built-in function `linprog` with the tolerance set to 10^{-8} has been used for solving all linear programs. In all experiments, we have used Algorithm 2 with the number of smallest eigenpairs included in \mathcal{V} set to $\ell = 1$, since this already provided significant improvements over Algorithm 1. For choosing r from Section 3.2, we have tested all values $r = 0, 1, \dots, Q$, see Remark 3.3.

Remark 4.1. *A slight modification of Algorithm 1 can significantly reduce the time spent on solving linear programs. For $\mu \in \Xi$, suppose that $y_{\text{LB}}(\mu) \in \mathcal{Y}_{\text{LB}}(C_J)$ is a minimizer of (7) on $\mathcal{Y}_{\text{LB}}(C_J)$. Let (λ_{J+1}, v_{J+1}) be the smallest eigenpair of $A(\mu_{J+1})$ computed in iteration $J + 1$. If $\theta(\mu)^T y_{\text{LB}}(\mu) \geq \lambda_{J+1}$, we have $y_{\text{LB}}(\mu) \in \mathcal{Y}_{\text{LB}}(C_{J+1}) \subseteq \mathcal{Y}_{\text{LB}}(C_J)$, making $y_{\text{LB}}(\mu)$ also a minimizer of (7) on $\mathcal{Y}_{\text{LB}}(C_{J+1})$. In this case, we have $\lambda_{\text{LB}}(\mu; C_{J+1}) = \lambda_{\text{LB}}(\mu; C_J)$ and there is no need to solve the linear program in (7).*

4.1 Random matrices

We first consider an academic example, where a random dense Hermitian matrix $A_1 \in \mathbb{C}^{N \times N}$ is perturbed, to a certain extent, by random Hermitian matrices $A_2, \dots, A_Q \in \mathbb{C}^{N \times N}$:

$$A(\mu) = A_1 + \mu_2 A_2 + \dots + \mu_Q A_Q,$$

where $\mu = (\mu_2, \dots, \mu_Q) \in D = [0, \delta]^{Q-1}$.

Example 4.2. *We consider $Q = 4$, $N = 1000$, $\delta = 0.2$, with A_1, A_2, A_3, A_4 having real random entries from the unit normal distribution. The performances of both algorithms is shown in Figure 3. The convergence of Algorithm 1 flattens after around 25 iterations and does not reach the desired tolerance, while the convergence of Algorithm 2 is much faster and reaches the desired tolerance within 47 iterations. We have also considered the optimized version of Algorithm 1, as described in Remark 4.1. To make the comparison fair, we have compared it to a variant of Algorithm 2 where the subspace bounds are recomputed only when $\lambda_{\text{LB}}(\mu; C_J)$ is recomputed. The influence of these modifications on the performances of both algorithms can be seen in Figure 4. The optimized version of Algorithm 2 requires a slightly larger number of iterations to converge to prescribed accuracy, but it still outperforms even the optimized version of Algorithm 1. Since Algorithm 2 converges quickly, there is no need to even consider the residual-based lower*

bounds from Section 3.5, but we still include the results in Figures 3 and 4 for the sake of completeness. Here and in the following, the star denotes the iteration from which on the residual-based lower bound (30) actually constitutes a lower bound for the smallest eigenvalue.




Figure 3: Convergence plots for Algorithms 1 and 2 applied to Example 4.2.

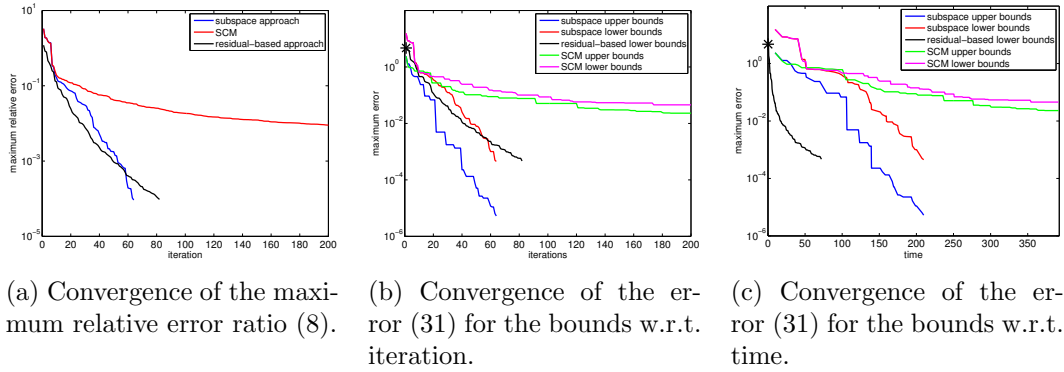


Figure 4: Convergence plots for the optimized versions of Algorithms 1 and 2 applied to Example 4.2.

4.2 Estimation of the coercivity constant

A posteriori error estimation in model order reduction techniques for parametrized PDEs, such as reduced basis method, requires reliable estimates for the coercivity constant [34] defined as

$$\alpha(\mu) = \inf_{u \in X} \frac{a(u, u; \mu)}{\|u\|_X^2}. \quad (32)$$

Here, $a(\cdot, \cdot, \mu)$ is the coercive symmetric bilinear form in the weak formulation of the underlying PDE on the domain Ω and X is the accompanying function space with the norm $\|\cdot\|_X$ induced by the scalar product $(u, v)_X = \tau(u, v)_{L^2(\Omega)} + a(u, v; \bar{\mu})$, with $\tau > 0$ and a fixed parameter value $\bar{\mu}$ chosen to be the center of D . A finite element discretization of (32) leads to the minimization problem

$$\alpha^N(\mu) = \inf_{v \in \mathbb{R}^N} \frac{v^T A(\mu) v}{v^T X v}, \quad (33)$$

where $A(\mu), X \in \mathbb{R}^{N \times N}$ are the matrices discretizing $a(\cdot, \cdot, \mu)$ and $(\cdot, \cdot)_X$, respectively. Minimizing (33) is clearly equivalent to computing the smallest eigenvalue of the generalized eigenvalue problem

$$A(\mu)v = \lambda Xv.$$

By computing a (sparse) Cholesky factorization $X = LL^T$, we transform (1):

$$L^{-1}A(\mu)L^{-T}w = \lambda w.$$

Hence, the matrices A_i appearing in Assumption 1.1 need to be replaced by

$$L^{-1}A_iL^{-T}, \quad i = 1, \dots, Q.$$

In the following, we consider three numerical examples of this type from the rbMIT toolbox [13]. We only include brief explanations of the examples; more details can be found in [13] and [30].


Example 4.3. *This example concerns a linear elasticity model of a parametrized body (see Figure 5a). The parameter μ_1 determines the width of the hole in the body while the parameter μ_2 determines its Poisson's ratio. A discretization of the underlying PDE leads to the matrix $A(\mu) = \sum_{i=1}^Q \theta_i(\mu)A_i$, with $Q = 16$, $\mu = (\mu_1, \mu_2)$ and functions $\theta_i(\mu)$ that arise from the parametrization of the geometry. We choose $N = 2183$ and $D = [-0.1, 0.1] \times [0.2, 0.3]$. As can be seen from Figure 5, The results are similar to those presented in Example 4.2, with Algorithm 2 converging in 31 iteration and Algorithm 1 not reaching the desired tolerance.*

Example 4.4. *This example concerns a stationary heat equation on a parametrized domain (see Figure 6a). The parameter μ_1 determines the coefficient in the Robin boundary conditions while the parameter μ_2 determines the length of the domain. A discretization of the underlying PDE leads to the matrix $A(\mu) = \sum_{i=1}^Q \theta_i(\mu)A_i$, with $Q = 3$, $\mu = (\mu_1, \mu_2)$ and functions $\theta_i(\mu)$ arising from the parametrization of the geometry and boundary conditions. We choose $N = 1311$ and $D = [0.02, 0.5] \times [2, 8]$. As can be seen from Figure 6, the results are similar to those observed in Examples 4.2 and 4.3.*


Example 4.5. *This example concerns a stationary heat equation on a square domain divided into blocks (see Figure 7a). In each of the subdomains, one of the parameters μ_1, \dots, μ_9 determines a coefficient of the PDE*

$$\operatorname{div} \left(\begin{bmatrix} 1 & -\mu_i \\ -\mu_i & 1 \end{bmatrix} \nabla u \right) = 0 \text{ on } \Omega_i, \quad i = 1, \dots, 9.$$

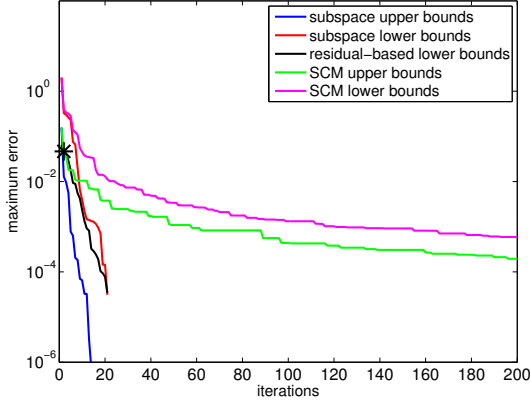
A discretization of the PDE leads to the matrix $A(\mu) = \sum_{i=1}^Q \theta_i(\mu)A_i$, where $Q = 10$, $\mu = (\mu_1, \dots, \mu_9)$ and functions $\theta_i(\mu)$ arising from the parametrization of the PDE coefficients. We choose $N = 1056$ and $D = [0.1, 0.5]^9$. As can be seen in Figure 7, the performance of both Algorithms 1 and 2 is not satisfactory, due to the slow convergence of the SCM and subspace lower bounds. Only the subspace upper bounds converges at a satisfactory rate. In this example, the residual-based lower bounds clearly show their advantage. They become reliable after only 31 iterations.




(a) Geometry of the underlying PDE.



(b) Convergence of the maximum relative error ratio (8).



(c) Convergence of the error (31) for the bounds w.r.t. iteration.



(d) Convergence of the error (31) for the bounds w.r.t. time.

Figure 5: Convergence plots for Algorithms 1 and 2 applied to Example 4.3.

5 Extension to computation of singular values

In Section 4.2 we have seen that the computation of coercivity constants can be formulated in terms of (1). For non-elliptic parametrized PDE one may have to resort to the inf-sup constant [11] defined as


$$\alpha(\mu) = \inf_{u \in X} \sup_{v \in X} \frac{a(u, v; \mu)}{\|u\|_X \|v\|_X}, \quad (34)$$

where $a(\cdot, \cdot, \mu)$ is the bilinear form in the weak formulation of the underlying PDE and X . A finite element discretization of (34) leads to the minimization problem

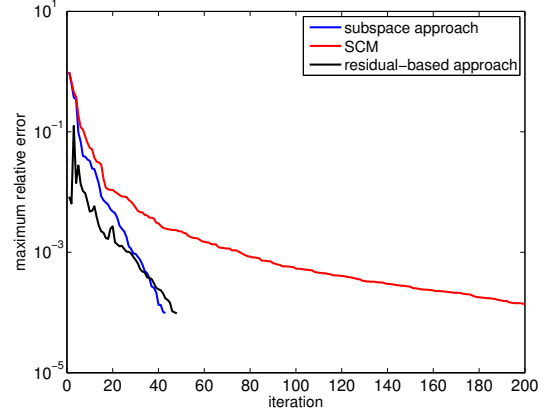
$$\inf_{u \in \mathbb{R}^N} \sup_{v \in \mathbb{R}^N} \frac{u^T A(\mu) v}{\sqrt{u^T X u} \sqrt{v^T X v}} = \inf_{x \in \mathbb{R}^N} \sup_{y \in \mathbb{R}^N} \frac{x^T L^{-T} A(\mu) L^{-1} y}{\|x\|_2 \|y\|_2} \quad (35)$$

where, once again, $A(\mu)$ and $X = LL^T$ are the discretizations of $a(\cdot, \cdot, \mu)$ and $(\cdot, \cdot)_X$, respectively. Minimizing (35) is equivalent to solving the singular value problem


$$\sigma_{\min}(L^{-1} A(\mu) L^{-T}),$$



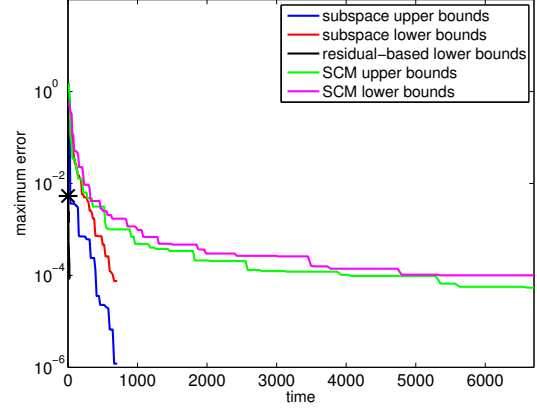
(a) Geometry of the underlying PDE.



(b) Convergence of the maximum relative error ratio (8).



(c) Convergence of the error (31) for the bounds w.r.t. iteration.



(d) Convergence of the error (31) for the bounds w.r.t. time.

Figure 6: Convergence plots for Algorithm 1 and 2 applied to Example 4.4.


which, in turn, is equivalent to computing

$$\lambda_{\min}(L^{-1}A(\mu)^T X^{-1}A(\mu)L^{-T}), \quad (36)$$

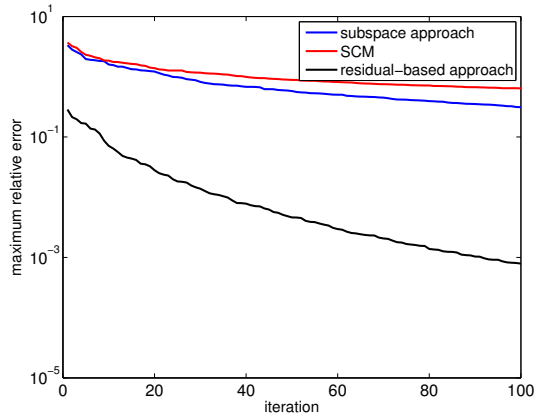
since $\sigma_{\min}(A) = \sqrt{\lambda_{\min}(A^T A)}$. The expression (36) can be recast in terms of (1), with Q^2 terms, with the matrices $A_{i,j}$ and functions $\theta_{ij}(\mu)$ for $i, j = 1, \dots, Q$ defined as

$$\begin{aligned} A_{ij} &= L^{-1}A_i^T X^{-1}A_j L^{-T} \\ \theta_{ij}(\mu) &= \theta_i(\mu)\theta_j(\mu). \end{aligned}$$

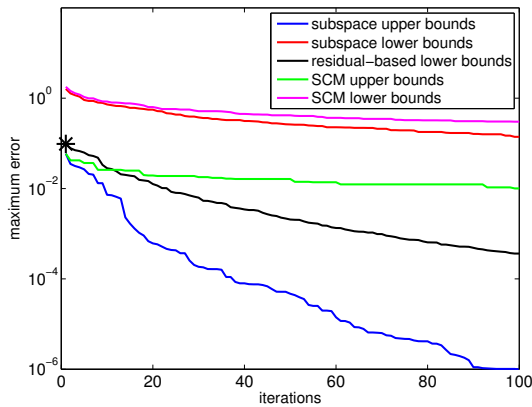
The SCM algorithm has already been applied to (36) but only with limited success, since having Q^2 terms in the affine decomposition of $A(\mu)$ further increases the computational cost by making the solution of the LP problem (7) significantly harder. The faster convergence of the subspace-accelerated approach to (36) mitigates this cost to a certain extent.




(a) Geometry of the underlying PDE.



(b) Convergence of the maximum relative error ratio (8).



(c) Convergence of the error (31) for the bounds w.r.t. iteration.



(d) Convergence of the error (31) for the bounds w.r.t. time.

Figure 7: Convergence plots for Algorithm 1 and 2 applied to Example 4.5.

6 Conclusions

Solving a parametrized Hermitian eigenvalue problem can be computationally very hard and SCM is the most commonly used existing approach. We have proposed a new subspace-accelerated approach, given in Algorithm 2. As can be seen in Section 3.4, it has better theoretical properties than SCM. As can be seen in Section 4, it also improves significantly on SCM in practice, for a number of examples discussed in the literature, while having only slightly larger computational cost per iteration. For problems with small gaps between the smallest eigenvalues, as in Example 4.5, the convergence of the subspace lower bounds may still not be satisfactory. For such cases, we propose a heuristic approach using residual-based lower bounds. The proposed approach can be extended to the solution of parametrized singular value problems.

7 Acknowledgements

We thank Christine Tobler and Meiyue Shao for discussing various ideas and approaches with us.

A Proof of Lemma 3.1

As a composition of continuous functions, the function f is clearly continuous. To prove monotonicity we distinguish two cases. First, let $\eta \geq \lambda_{\mathcal{V}}^{(1)}$. Then

$$f(\eta) = \lambda_{\mathcal{V}}^{(1)} - 2\rho^2 / \left(\eta - \lambda_{\mathcal{V}}^{(1)} + \sqrt{(\eta - \lambda_{\mathcal{V}}^{(1)})^2 + 4\rho^2} \right),$$

which clearly increases as η increases. Now, let $\eta \leq \lambda_{\mathcal{V}}^{(1)}$. Then

$$f(\eta) = \eta - 2\rho^2 / \left(\lambda_{\mathcal{V}}^{(1)} - \eta + \sqrt{(\eta - \lambda_{\mathcal{V}}^{(1)})^2 + 4\rho^2} \right)$$

and

$$f'(\eta) = 1 - \frac{2\rho^2}{\left(\lambda_{\mathcal{V}}^{(1)} - \eta + \sqrt{(\lambda_{\mathcal{V}}^{(1)} - \eta)^2 + 4\rho^2} \right) \sqrt{(\lambda_{\mathcal{V}}^{(1)} - \eta)^2 + 4\rho^2}}.$$

Showing $f'(\eta) \geq 0$, and thus establishing monotonicity, is equivalent to

$$\begin{aligned} (\lambda_{\mathcal{V}}^{(1)} - \eta)^2 + 4\rho^2 + (\lambda_{\mathcal{V}}^{(1)} - \eta) \sqrt{(\lambda_{\mathcal{V}}^{(1)} - \eta)^2 + 4\rho^2} &\geq 2\rho^2 \\ (\lambda_{\mathcal{V}}^{(1)} - \eta) \sqrt{(\lambda_{\mathcal{V}}^{(1)} - \eta)^2 + 4\rho^2} &\geq 0 \geq -(\lambda_{\mathcal{V}}^{(1)} - \eta)^2 - 2\rho^2, \end{aligned}$$

which is trivially satisfied for $\lambda_{\mathcal{V}}^{(1)} \geq \eta$. This completes the proof.

References

- [1] R. Andreev and C. Schwab. Sparse tensor approximation of parametric eigenvalue problems. In *Numerical analysis of multiscale problems*, volume 83 of *Lect. Notes Comput. Sci. Eng.*, pages 203–241. Springer, Heidelberg, 2012.
- [2] Z. Bai, J. W. Demmel, J. J. Dongarra, A. Ruhe, and H. van der Vorst, editors. *Templates for the solution of algebraic eigenvalue problems*, volume 11 of *Software, Environments, and Tools*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2000. A practical guide.
- [3] 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. *C. R. Math. Acad. Sci. Paris*, 339(9):667–672, 2004.
- [4] T. Betcke and L. N. Trefethen. Reviving the method of particular solutions. *SIAM Rev.*, 47(3):469–491, 2005.
- [5] J. De Vlieger and K. Meerbergen. A subspace method for unimodal symmetric eigenvalue optimization problems involving large scale matrices. *TW Reports*, 2012.

- [6] C. Engström, C. Hafner, and K. Schmidt. Computation of lossy Bloch waves in two-dimensional photonic crystals. *J. Comput. Theor. Nanosci.*, 6:1–9, 2009.
- [7] A. Ghosh, S. Boyd, and A. Saberi. Minimizing effective resistance of a graph. *SIAM Rev.*, 50(1):37–66, 2008.
- [8] E. Gutkin, E. A. Jonckheere, and M. Karow. Convexity of the joint numerical range: topological and differential geometric viewpoints. *Linear Algebra Appl.*, 376:143–171, 2004.
- [9] C. Helmberg and F. Rendl. A spectral bundle method for semidefinite programming. *SIAM J. Optim.*, 10(3):673–696 (electronic), 2000.
- [10] R. A. Horn and C. R. Johnson. *Matrix analysis*. Cambridge University Press, Cambridge, second edition, 2013.
- [11] D. B. P. Huynh, D. J. Knezevic, Y. Chen, J. S. Hesthaven, and A. T. Patera. A natural-norm successive constraint method for inf-sup lower bounds. *Comput. Methods Appl. Mech. Engrg.*, 199(29-32):1963–1975, 2010.
- [12] D. B. P. Huynh, G. Rozza, S. Sen, and A. T. Patera. 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.
- [13] D.B.P. Huynh, Nguyen N.C., A.T. Patera, and G. Rozza. rbmit software, 2010.
- [14] C. R. Johnson. A Gersgorin-type lower bound for the smallest singular value. *Linear Algebra Appl.*, 112:1–7, 1989.
- [15] T. Kato. *Perturbation theory for linear operators*. Classics in Mathematics. Springer-Verlag, Berlin, 1995. Reprint of the 1980 edition.
- [16] Y. Kim and M. Mesbahi. On maximizing the second smallest eigenvalue of a state-dependent graph Laplacian. *IEEE Trans. Autom. Control*, 51(1):116–120, 2006.
- [17] A. V. Knyazev. Toward the optimal preconditioned eigensolver: Locally optimal block preconditioned conjugate gradient method. *SIAM J. Sci. Comput.*, 23(2):517–541, 2001.
- [18] D. Kressner and C. Tobler. Low-rank tensor Krylov subspace methods for parametrized linear systems. *SIAM J. Matrix Anal. Appl.*, 32(4):1288–1316, 2011.
- [19] D. Kressner and B. Vandereycken. Subspace methods for computing the pseudospectral abscissa and the stability radius. *SIAM J. Matrix Anal. Appl.*, 35(1):292–313, 2014.
- [20] R. B. Lehoucq, D. C. Sorensen, and C. Yang. *ARPACK users’ guide*. SIAM, Philadelphia, PA, 1998. Solution of large-scale eigenvalue problems with implicitly restarted Arnoldi methods.
- [21] A. S. Lewis and M. L. Overton. Eigenvalue optimization. In *Acta numerica, 1996*, volume 5 of *Acta Numer.*, pages 149–190. Cambridge Univ. Press, Cambridge, 1996.

- [22] C.-K. Li and R.-C. Li. A note on eigenvalues of perturbed Hermitian matrices. *Linear Algebra Appl.*, 395:183–190, 2005.
- [23] L. Machiels, Y. Maday, I. B. Oliveira, A. T. Patera, and D. V. Rovas. Output bounds for reduced-basis approximations of symmetric positive definite eigenvalue problems. *C. R. Math. Acad. Sci. Paris*, 331(2):153–158, 2000.
- [24] A. Manzoni and F. Negri. Rigorous and heuristic strategies for the approximation of stability factors in nonlinear parametrized PDEs. Technical report MATHICSE 8.2014, 2014.
- [25] J. C. Mason and D. C. Handscomb. *Chebyshev polynomials*. Chapman & Hall/CRC, Boca Raton, FL, 2003.
- [26] J. Matousek and B. Gärtner. *Understanding and using linear programming*. Springer Science & Business Media, 2007.
- [27] E. Mengi, E. A. Yildirim, and M. Kiliç. Numerical optimization of eigenvalues of Hermitian matrix functions. *SIAM J. Matrix Anal. Appl.*, 35(2):699–724, 2014.
- [28] N. C. Nguyen, K. Veroy, and A. T. Patera. Certified real-time solution of parametrized partial differential equations. In *Handbook of Materials Modeling*, pages 1529–1564. Springer, 2005.
- [29] B. N. Parlett. *The Symmetric Eigenvalue Problem*, volume 20 of *Classics in Applied Mathematics*. SIAM, Philadelphia, PA, 1998. Corrected reprint of the 1980 original.
- [30] A.T. Patera and G. Rozza. *Reduced Basis Approximation and A Posteriori Error Estimation for Parametrized Partial Differential Equations*. MIT-Pappalardo Graduate Monographs in Mechanical Engineering. Cambridge, MA, US, 2012. To appear. Preliminary version available from <http://augustine.mit.edu/methodology/>.
- [31] C. Prud’homme, D. V. Rovas, K. Veroy, L. Machiels, Y. Maday, A. T. Patera, and G. Turinici. A mathematical and computational framework for reliable real-time solution of parametrized partial differential equations. *M2AN Math. Model. Numer. Anal.*, 36(5):747–771, 2002.
- [32] M. Reed and B. Simon. *Methods of modern mathematical physics. IV. Analysis of operators*. Academic Press [Harcourt Brace Jovanovich Publishers], New York, 1978.
- [33] M. Rojas, S. A. Santos, and D. C. Sorensen. A new matrix-free algorithm for the large-scale trust-region subproblem. *SIAM J. Optim.*, 11(3):611–646, 2000/01.
- [34] G. Rozza, D.B.P. Huynh, and A. T. Patera. Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations: application to transport and continuum mechanics. *Arch. Comput. Methods Eng.*, 15(3):229–275, 2008.
- [35] S. Sen, K. Veroy, D. B. P. Huynh, S. Deparis, N. C. Nguyen, and A. T. Patera. “Natural norm” a posteriori error estimators for reduced basis approximations. *J. Comput. Phys.*, 217(1):37–62, 2006.

- [36] G. W. Stewart and J.-G. Sun. *Matrix Perturbation Theory*. Academic Press, New York, 1990.
- [37] L. N. Trefethen and M. Embree. *Spectra and pseudospectra*. Princeton University Press, Princeton, NJ, 2005. The behavior of nonnormal matrices and operators.
- [38] K. Veroy, D. V. Rovas, and A. T. Patera. A posteriori error estimation for reduced-basis approximation of parametrized elliptic coercive partial differential equations: “convex inverse” bound conditioners. *ESAIM Control Optim. Calc. Var.*, 8:1007–1028, 2002.

Recent publications:

**MATHEMATICS INSTITUTE OF COMPUTATIONAL SCIENCE AND ENGINEERING
Section of Mathematics
Ecole Polytechnique Fédérale
CH-1015 Lausanne**

- 46.2014** ASSYR ABDULLE:
The role of numerical integration in numerical in numerical homogenization
- 47.2014** ANDREA MANZONI, STEFANO PAGANI, TONI LASSILA:
Accurate solution of Bayesian inverse uncertainty quantification problems using model and error reduction methods
- 48.2014** MARCO PICASSO:
From the free surface flow of a viscoelastic fluid towards the elastic deformation of a solid
- 49.2014** FABIO NOBILE, FRANCESCO TESEI:
A multi level Monte Carlo method with control variate for elliptic PDEs with log-normal coefficients
- ***
- 01.2015** PENG CHEN, ALFIO QUARTERONI, GIANLUIGI ROZZA:
Reduced order methods for uncertainty quantification problems
- 02.2015** FEDERICO NEGRI, ANDREA MANZONI, DAVID AMSALLEM:
Efficient model reduction of parametrized systems by matrix discrete empirical interpolation
- 03.2015** GIOVANNI MIGLIORATI, FABIO NOBILE, RAÚL TEMPONE:
Convergence estimate in probability and in expectation for discrete least squares with noisy evaluations at random points
- 04.2015** FABIO NOBILE, LORENZO TAMELLINI, FRANCESCO TESEI, RAÚL TEMPONE:
An adaptive sparse grid algorithm for elliptic PDEs with lognormal diffusion coefficient
- 05.2015** MICHAEL STEINLECHNER:
Riemannian optimization for high-dimensional tensor completion
- 06.2015** V. R. KOSTIĆ, A. MIEDLAR, LJ. CVETKOVIĆ:
An algorithm for computing minimal Geršgorin sets
- 07.2015** ANDREA BARTEZZAGHI, LUCA DEDÈ, ALFIO QUARTERONI:
Isogeometric analysis of high order partial differential equations on surfaces
- 08.2015** IVAN FUMAGALLI, ANDREA MANZONI, NICOLA PAROLINI, MARCO VERANI:
Reduced basis approximation and a posteriori error estimates for parametrized elliptic eigenvalue problems
- 09.2015** DAVIDE FORTI, LUCA DEDÈ:
Semi-implicit BDF time discretization of the Navier-Stokes equations with VMS-LES modeling in a High Performance Computing framework
- 10.2015** PETAR SIRKOVIĆ, DANIEL KRESSNER:
Subspace acceleration for large-scale parameter-dependent Hermitian eigenproblems