

Randomized linear algebra for model reduction. Part II: minimal residual methods and dictionary-based approximation.

Oleg Balabanov^{*†} and Anthony Nouy^{*‡}

Abstract

A methodology for using random sketching in the context of model order reduction for high-dimensional parameter-dependent systems of equations was introduced in [Balabanov and Nouy 2019, Part I]. Following this framework, we here construct a reduced model from a small, efficiently computable random object called a sketch of a reduced model, using minimal residual methods. We introduce a sketched version of the minimal residual based projection as well as a novel nonlinear approximation method, where for each parameter value, the solution is approximated by minimal residual projection onto a subspace spanned by several vectors picked (online) from a dictionary of candidate basis vectors. It is shown that random sketching technique can improve not only efficiency but also numerical stability. A rigorous analysis of the conditions on the random sketch required to obtain a given accuracy is presented. These conditions may be ensured a priori with high probability by considering for the sketching matrix an oblivious embedding of sufficiently large size. Furthermore, a simple and reliable procedure for a posteriori verification of the quality of the sketch is provided. This approach can be used for certification of the approximation as well as for adaptive selection of the size of the random sketching matrix. We also propose a two-step procedure for an efficient and stable estimation of an inner product between parameter-dependent vectors having affine decompositions with many terms (possibly expensive to maintain). This procedure can be used for extraction of a quantity of interest (linear or quadratic functional of the solution), estimation of the primal-dual correction, etc.

Keywords— model reduction, reduced basis, random sketching, subspace embedding, minimal residual methods, sparse approximation, dictionary

1 Introduction

We consider large parameter-dependent systems of equations

$$\mathbf{A}(\mu)\mathbf{u}(\mu) = \mathbf{b}(\mu), \quad \mu \in \mathcal{P}, \quad (1)$$

^{*}Centrale Nantes, LMJL, UMR CNRS 6629, France.

[†]Polytechnic University of Catalonia, LaCàn, Spain.

[‡]Corresponding author (anthony.nouy@ec-nantes.fr).

where $\mathbf{u}(\mu)$ is a solution vector, $\mathbf{A}(\mu)$ is a parameter-dependent matrix, $\mathbf{b}(\mu)$ is a parameter-dependent right hand side and \mathcal{P} is a parameter set. Parameter-dependent problems are considered for many purposes such as design, control, optimization, uncertainty quantification or inverse problems.

Solving (1) for many parameter values can be computationally unfeasible. Moreover, for real-time applications, a quantity of interest ($\mathbf{u}(\mu)$ or a function of $\mathbf{u}(\mu)$) has to be estimated on the fly in highly limited computational time for a certain value of μ . Model order reduction (MOR) methods are developed for efficient approximation of the quantity of interest for each parameter value. They typically consist of two stages. In the first so-called offline stage a reduced model is constructed from the full order model. This stage usually involves expensive computations such as evaluations of $\mathbf{u}(\mu)$ for several parameter values, computing multiple high-dimensional matrix-vector and inner products, etc., but this stage is performed only once. Then, for each given parameter value, the precomputed reduced model is used for efficient approximation of the solution or an output quantity with a computational cost independent of the dimension of the initial system of equations (1). For a detailed presentation of the classical MOR methods such as Reduced Basis (RB) method and Proper Orthogonal Decomposition (POD) the reader can refer to [4]. In the present work the approximation of the solution shall be obtained with a minimal residual (minres) projection on a reduced (possibly parameter-dependent) subspace. The minres projection can be interpreted as a Petrov-Galerkin projection where the test space is chosen to minimize some norm of the residual [1, 9]. Major benefits over the classical Galerkin projection include an improved stability (quasi-optimality) for non-coercive problems and more effective residual-based error bounds of an approximation (see e.g. [9]). In addition, minres methods are better suited to random sketching as will be seen in the present article.

In recent years randomized linear algebra (RLA) became a popular approach in the fields such as data analysis, machine learning, compressed sensing, etc. [27, 32, 33]. This probabilistic approach for numerical linear algebra can yield a drastic computational cost reduction in terms of classical metrics of efficiency such as complexity (number of flops) and memory consumption. Moreover, it can be highly beneficial in extreme computational environments that are typical in contemporary scientific computing. For instance, RLA can be essential when data has to be analyzed only in one pass (e.g., when it is streamed from a server) or when it is distributed on multiple workstations with expensive communication costs.

Despite their indisputable success in fields closely related to MOR, the aforementioned techniques only recently started to be used for model order reduction. One of the earliest works considering RLA in the context of MOR is [34], where the authors proposed to use RLA for interpolation of (implicit) inverse of a parameter-dependent matrix. In [8] the RLA was used for approximating the range of a transfer operator and for computing a probabilistic bound for the approximation error. In [20, 30] the authors developed probabilistic error estimators based on random projections and the adjoint method. These approaches can also be formulated in RLA framework. Efficient algorithms for Dynamic Mode Decomposition based on RLA and compressed sensing were proposed in [6, 7, 16, 17, 24], though with a lack of theoretical analysis.

As already shown in [3], random sketching can lead to drastic reduction of the computational costs of classical MOR methods. A random sketch of a reduced model is defined as a set of small

random projections of the reduced basis vectors and the associated residuals. Its representation (i.e., affine decomposition¹) can be efficiently precomputed in basically any computational architecture. The random projections should be chosen according to the metric of efficiency, e.g., number of flops, memory consumption, communication cost between distributed machines, scalability, etc. A rigorous analysis of the cost of obtaining a random sketch in different computational environments can be found in [3, Section 4.4]. When a sketch has been computed, the reduced model can be approximated without operating on large vectors but only on their small sketches typically requiring negligible computational costs. The approximation can be guaranteed to almost preserve the quality of the original reduced model with user-specified probability. The computational cost depends only logarithmically on the probability of failure, which can therefore be chosen very small, say 10^{-10} . In [3] it was shown how random sketching can be employed for an efficient approximation of the Galerkin projection, the computation of the norm of the residual for error estimation, and the computation of a primal-dual correction. Furthermore, new efficient sketched versions of the weak greedy algorithm and Proper Orthogonal Decomposition were introduced for generation of reduced bases.

1.1 Contributions

The present work is a continuation of [3], where we adapt the random sketching technique to minimal residual methods, propose a dictionary-based approximation method and additionally discuss the questions of a posteriori certification of the sketch and efficient extraction of the quantity of interest from a solution of the reduced model. A detailed discussion on the major contributions of the paper is provided below.

First a sketched version of the minres projection is proposed in Section 3, which is more efficient (in both offline and online stages) and numerically stable than the classical approach. The classical construction of the reduced order model with minres projection involves the evaluation of multiple inner products, which can become a burden for high-dimensional problems. Furthermore, the classical procedure for ensuring stability of the reduced (normal) system of equations² only guarantees that the condition number is bounded by the square of the condition number of $\mathbf{A}(\mu)$. Such a bound can be insufficient for applications with very or even moderately ill-conditioned operators. In addition, the online evaluation of the reduced system of equations from the precomputed affine expansions can also be very expensive and suffer from round-off errors. The random sketching technique can offer more efficient (in both offline and online stages) and numerically stable procedures for estimating solutions to the minimal residual problem. Here the reduced model is approximated from its efficiently computable random sketch with a negligible computational cost. The precomputation of a random sketch can require much lower complexity, memory consumption and communication cost than the computations involved in the classical offline stage. As shown in Section 3.2, with random sketching the online solution can be found by solving a small $\mathcal{O}(r)$ by r least-squares problem. The construction of this problem takes only $\mathcal{O}(r^2 m_A + r m_b)$ flops

¹A parameter-dependent quantity $\mathbf{v}(\mu)$ with values in a vector space V over \mathbb{K} is said to admit an affine representation if $\mathbf{v}(\mu) = \sum \mathbf{v}_i \lambda_i(\mu)$ with $\lambda_i(\mu) \in \mathbb{K}$ and $\mathbf{v}_i \in V$.

²This procedure consists in orthogonalization of the basis corresponding to some inner product.

(compared to $\mathcal{O}(r^2 m_A^2 + r m_b^2)$ flops required for forming the classical minres reduced system of equations), where m_A and m_b are the numbers of terms in the affine expansions of $\mathbf{A}(\mu)$ and $\mathbf{b}(\mu)$. Moreover, when the basis is orthogonalized one can guarantee a better stability of the reduced least-squares matrix (the condition number is bounded by the condition number of $\mathbf{A}(\mu)$). In addition, the parameter-dependent reduced matrix can have an affine expansion with considerably less terms and therefore its evaluation is less sensitive to round-off errors. It is also shown that the size of the sketching matrix which is sufficient to preserve the quasi-optimality constants of minres projection can be characterized regardless of the properties of the operator (e.g., the condition number). This feature proves the sketched minres projection to be more robust than the sketched Galerkin projection for which the preservation of the approximation's quality can degrade dramatically for ill-conditioned or non-coercive problems as was revealed in [3].

In Section 4 we introduce a novel nonlinear method for approximating the solution to (1), where for each value of the parameter, the solution is approximated by a minimal residual projection onto a subspace spanned by several vectors from a dictionary. From an algebraic point of view, this approach can be formulated as a parameter-dependent sparse least-squares problem. It is shown that the solution can be accurately estimated (with high probability) from a random sketch associated with the dictionary, which allows drastic reduction of the computational costs. A condition on the dimension of the random sketch required to obtain a given accuracy is provided. Again, the construction of a reduced model does not require operations on high-dimensional vectors but only on their sketches. In particular, in the offline stage we only need to maintain a sketch of the dictionary. In Section 4.4 we propose an efficient greedy procedure for the dictionary construction based on snapshots.

The dictionary-based approach is more natural than the classical *hp*-refinement method [14, 15] and it should always provide a better approximation (see Section 4.1). The potential of approximation with dictionaries for problems with a slow decay of the Kolmogorov r -widths of the solution manifold was revealed in [13, 23]. Although they improved classical approaches, the algorithms proposed in [13, 23] still involve in general heavy computations in both offline and online stages, and can suffer from round-off errors. If n and r are the dimensions of the full solution space and the (parameter-dependent) reduced approximation space, respectively, and K is the cardinality of the dictionary, then the offline complexity and the memory consumption associated with post-processing of the snapshots in [13, 23] are at least $\mathcal{O}(n(K^2 m_A^2 + m_b^2))$ and $\mathcal{O}(nK + K^2 m_A^2 + m_b^2)$, respectively. Furthermore, the online stage in [23] requires $\mathcal{O}((r^3 + m_A^2 r)K + m_b^2)$ flops and $\mathcal{O}(m_A^2 K^2 + m_b^2)$ bytes of memory. The high offline cost and the high cost of maintenance of the reduced model (which are proportional to K^2) limits the effectiveness of the method in many computational environments. Moreover, we see that the online complexity of the approach in [23] is proportional to Kr^3 , which leads to high computational costs for moderate r . Random sketching can drastically improve efficiency and stability of the dictionary-based approximation. The offline complexity and memory requirements of the construction of a reduced model with random sketching (see Section 4.4) are $\mathcal{O}(n(Km_A + m_b)(\log r + \log \log K))$ and $\mathcal{O}(n + (Km_A + m_b)r \log K)$, respectively. We observe reduction (compared to [13, 23]) of the complexity and memory consumption of the offline stage by at least a factor of $\mathcal{O}(K)$. In its turn, the online stage with random sketching needs $\mathcal{O}((m_A K + m_b)r \log K + r^2 K \log K)$ flops and $\mathcal{O}((m_A K + m_b)r \log K)$

bytes of memory, which are in $\mathcal{O}((r + m_A + \frac{m_b}{r})/\log K)$ and $\mathcal{O}((m_A K + m_b)/(r \log K))$ times less than the requirements in [23]. Note that in some places logarithmic terms were neglected. A more detailed analysis of the computational costs can be found in Sections 4.3 and 4.4.

The online stage usually proceeds with computation of the coordinates of an approximation $\mathbf{u}_r(\mu)$ of $\mathbf{u}(\mu)$ in a certain basis. Then the coordinates are used for the efficient evaluation of an estimation $s_r(\mu) := l(\mathbf{u}_r(\mu); \mu)$ of a quantity of interest $s(\mu) := l(\mathbf{u}(\mu); \mu)$ from an affine decomposition of $l(\cdot, \mu)$. When the affine decomposition of $l(\cdot, \mu)$ is expensive to maintain and to operate with, precomputation of the affine decomposition of $s_r(\mu)$ can become too cumbersome. This is the case when $l(\cdot, \mu)$ contains numerous terms in its affine decomposition or when one considers too large, possibly distributed, basis for $\mathbf{u}_r(\mu)$. In Section 5 we provide a way to efficiently estimate $s_r(\mu)$. Our procedure is two-phased. First the solution $\mathbf{u}_r(\mu)$ is approximated by a projection $\mathbf{w}_p(\mu)$ on a new basis, which is cheap to operate with. The affine decomposition of an approximation $s_r(\mu) \approx l(\mathbf{w}_p(\mu); \mu)$ can now be efficiently precomputed. Then in the second step, the accuracy of $l(\mathbf{w}_p(\mu); \mu)$ is improved with a random correction computable from the sketches of the two bases with a negligible computational cost. Note that our approach can be employed for the efficient computation of the affine decomposition of primal-dual corrections and quadratic quantities of interest (see Remark 5.1).

As shown in [3] for Galerkin methods and in Sections 3.2 and 4.3 of the present paper for minres methods, a sketch of a reduced model almost preserves the quality of approximation when the sketching matrix satisfies an ε -embedding property. Such a matrix may be generated randomly by considering an oblivious subspace embedding of sufficiently large size. The number of rows for the oblivious embedding may be selected with the theoretical bounds provided in [3]. However, it was revealed that these bounds are pessimistic or even impractical (e.g., for adaptive algorithms or POD). In practice, one can consider embeddings of much smaller sizes and still obtain accurate reduced order models. Moreover, for some random matrices, theoretical bounds may not be available although there might exist strong empirical evidence that these matrices should yield outstanding performances (e.g., matrices constructed with random Givens rotations as in [28]). When no a priori guarantee on the accuracy of the given sketching matrix is available or when conditions on the size of the sketch based on a priori analysis are too pessimistic, one can provide a posteriori guarantee. An easy and robust procedure for a posteriori verification of the quality of a sketch of a reduced model is provided in Section 6. The methodology can also be used for deriving a criterion for adaptive selection of the size of the random sketching matrix to yield an accurate estimation of the reduced model with high probability.

The outline of the article is as follows. In Section 2 we introduce the problem setting and recall the main ingredients of the framework developed in [3]. The minimal residual method considering a projection on a single low-dimensional subspace is presented in Section 3. We present a standard minres projection in a discrete form followed by its efficient approximation with random sketching. Section 4 presents a novel dictionary-based minimal residual method using random sketching. A two-phased procedure for efficient and stable extraction of the output quantity of interest from the reduced model's solution is proposed in Section 5. A posteriori verification of the quality of a sketch and few scenarios where such a procedure can be used are provided in Section 6. The methodology is validated numerically on two benchmark problems in Section 7.

2 Preliminaries

Let $\mathbb{K} = \mathbb{R}$ or \mathbb{C} and let $U := \mathbb{K}^n$ and $U' := \mathbb{K}^n$ represent the solution space and its dual, respectively. The solution $\mathbf{u}(\mu)$ is an element from U , $\mathbf{A}(\mu)$ is a linear operator from U to U' , the right hand side $\mathbf{b}(\mu)$ and the extractor of the quantity of interest $\mathbf{l}(\mu)$ are elements of U' .

Spaces U and U' are equipped with inner products $\langle \cdot, \cdot \rangle_U := \langle \mathbf{R}_U \cdot, \cdot \rangle$ and $\langle \cdot, \cdot \rangle_{U'} := \langle \cdot, \mathbf{R}_U^{-1} \cdot \rangle$, where $\langle \cdot, \cdot \rangle$ is the canonical ℓ_2 -inner product on \mathbb{K}^n and $\mathbf{R}_U : U \rightarrow U'$ is some symmetric (for $\mathbb{K} = \mathbb{R}$), or Hermitian (for $\mathbb{K} = \mathbb{C}$), positive definite operator. We denote by $\|\cdot\|$ the canonical ℓ_2 -norm on \mathbb{K}^n . Finally, for a matrix \mathbf{M} we denote by \mathbf{M}^H its (Hermitian) transpose.

2.1 Random sketching

A framework for using random sketching (see [19, 33]) in the context of MOR was introduced in [3]. The sketching technique is seen as a modification of the inner product in a given subspace (or a collection of subspaces). The modified inner product is an estimation of the original one and is much easier and more efficient to operate with. Next, we briefly recall the basic preliminaries from [3].

Let V be a subspace of U . The dual of V is identified with a subspace $V' := \{\mathbf{R}_U \mathbf{x} : \mathbf{x} \in V\}$ of U' . For a matrix $\Theta \in \mathbb{K}^{k \times n}$ with $k \leq n$ we define the following semi-inner products on U :

$$\langle \cdot, \cdot \rangle_U^\Theta := \langle \Theta \cdot, \Theta \cdot \rangle, \text{ and } \langle \cdot, \cdot \rangle_{U'}^\Theta := \langle \Theta \mathbf{R}_U^{-1} \cdot, \Theta \mathbf{R}_U^{-1} \cdot \rangle, \quad (2)$$

and we let $\|\cdot\|_U^\Theta$ and $\|\cdot\|_{U'}^\Theta$ denote the associated semi-norms.

Remark 2.1. *The extension of the methodology to the case where $\langle \cdot, \cdot \rangle_U$ is not definite, i.e., \mathbf{R}_U is positive semi-definite, is straightforward. Let us assume that $\langle \cdot, \cdot \rangle_U$ is an inner product on a subspace $W \subseteq U$ of interest. Then, it follows that $W' := \{\mathbf{R}_U \mathbf{x} : \mathbf{x} \in W\}$ can be equipped with $\langle \cdot, \cdot \rangle_{U'} := \langle \cdot, \mathbf{R}_U^\dagger \cdot \rangle$, where \mathbf{R}_U^\dagger is the pseudo-inverse of \mathbf{R}_U . Such products $\langle \cdot, \cdot \rangle_U$ and $\langle \cdot, \cdot \rangle_{U'}$ can be approximated by*

$$\langle \cdot, \cdot \rangle_U^\Theta := \langle \Theta \cdot, \Theta \cdot \rangle, \text{ and } \langle \cdot, \cdot \rangle_{U'}^\Theta := \langle \Theta \mathbf{R}_U^\dagger \cdot, \Theta \mathbf{R}_U^\dagger \cdot \rangle. \quad (3)$$

This will be useful for the estimation of a (semi-)inner product between parameter-dependent vectors in Section 5 (see Remark 5.2).

Definition 2.2. *A matrix Θ is called a $U \rightarrow \ell_2$ ε -subspace embedding (or simply an ε -embedding) for V , if it satisfies*

$$\forall \mathbf{x}, \mathbf{y} \in V, \quad |\langle \mathbf{x}, \mathbf{y} \rangle_U - \langle \mathbf{x}, \mathbf{y} \rangle_U^\Theta| \leq \varepsilon \|\mathbf{x}\|_U \|\mathbf{y}\|_U. \quad (4)$$

Here ε -embeddings shall be constructed as realizations of random matrices that are built in an oblivious way without any a priori knowledge of V .

Definition 2.3. *A random matrix Θ is called a (ε, δ, d) oblivious $U \rightarrow \ell_2$ subspace embedding if it is an ε -embedding for an arbitrary d -dimensional subspace $V \subset U$ with probability at least $1 - \delta$.*

Oblivious $\ell_2 \rightarrow \ell_2$ subspace embeddings (defined by Definition 2.2 with $\langle \cdot, \cdot \rangle_U := \langle \cdot, \cdot \rangle$) include the rescaled Gaussian distribution, the rescaled Rademacher distribution, the Subsampled Randomized Hadamard Transform (SRHT), the Subsampled Randomized Fourier Transform (SRFT), CountSketch matrix, SRFT combined with sequences of random Givens rotations, and others [3, 19, 28, 33]. In this work we shall rely on the rescaled Gaussian distribution and SRHT.

An oblivious $U \rightarrow \ell_2$ subspace embedding for a general inner product $\langle \cdot, \cdot \rangle_U$ can be constructed as

$$\Theta = \Omega \mathbf{Q}, \quad (5)$$

where Ω is a $\ell_2 \rightarrow \ell_2$ subspace embedding and $\mathbf{Q} \in \mathbb{K}^{s \times n}$ is an easily computable (possibly rectangular) matrix such that $\mathbf{Q}^H \mathbf{Q} = \mathbf{R}_U$ (see [3, Remark 2.7]).

It follows that an $U \rightarrow \ell_2$ ε -subspace embedding for V can be obtained with high probability as a realization of an oblivious subspace embedding of sufficiently large size.

The a priori estimates for the required size of Θ are usually pessimistic for practical use. Moreover, a good performance of certain random embeddings (e.g., matrices with sequences of random Givens rotations) was validated only empirically [19]. Therefore, in Section 6 we provide a simple and reliable a posteriori procedure for characterizing the quality of an embedding for each given subspace. Such a procedure can be used for the adaptive selection of the number of rows for Θ or for certifying the quality of the sketched reduced order model.

2.2 A sketch of a reduced model

Here the output of a reduced order model is efficiently estimated from its random sketch. The Θ -sketch of a reduced model associated with a subspace U_r is defined as

$$\{ \{ \Theta \mathbf{x}, \Theta \mathbf{R}_U^{-1} \mathbf{r}(\mathbf{x}; \mu) \} : \mathbf{x} \in U_r \}, \quad (6)$$

where $\mathbf{r}(\mathbf{x}; \mu) := \mathbf{b}(\mu) - \mathbf{A}(\mu) \mathbf{x}$. Let $\mathbf{U}_r \in \mathbb{K}^{n \times r}$ be a matrix whose columns form a basis of U_r . Then each element of (6) can be characterized from the coordinates of \mathbf{x} associated with \mathbf{U}_r , i.e., a vector $\mathbf{a}_r \in \mathbb{K}^r$ such that $\mathbf{x} = \mathbf{U}_r \mathbf{a}_r$, and the following quantities

$$\mathbf{U}_r^\Theta := \Theta \mathbf{U}_r, \quad \mathbf{V}_r^\Theta(\mu) := \Theta \mathbf{R}_U^{-1} \mathbf{A}(\mu) \mathbf{U}_r \quad \text{and} \quad \mathbf{b}^\Theta(\mu) := \Theta \mathbf{R}_U^{-1} \mathbf{b}(\mu). \quad (7)$$

Clearly $\mathbf{V}_r^\Theta(\mu)$ and $\mathbf{b}^\Theta(\mu)$ have affine expansions containing at most as many terms as the ones of $\mathbf{A}(\mu)$ and $\mathbf{b}(\mu)$, respectively. The matrix \mathbf{U}_r^Θ and the affine expansions of $\mathbf{V}_r^\Theta(\mu)$ and $\mathbf{b}^\Theta(\mu)$ are referred to as the Θ -sketch of \mathbf{U}_r (a representation of the Θ -sketch of a reduced model associated with U_r). With a good choice of an oblivious embedding, a Θ -sketch of \mathbf{U}_r can be efficiently precomputed in any computational environment (see [3]). Thereafter, an approximation of a reduced order model can be obtained with a negligible computational cost. Note that in [3] the affine expansion of $\mathbf{l}_r(\mu)^H := \mathbf{U}_r^H \mathbf{l}(\mu)$, where $\mathbf{l}(\mu) \in U'$ is an extractor of the linear quantity of interest, are also considered as a part of the Θ -sketch of \mathbf{U}_r . In the present paper, however, we consider a more general scenario where the computation of the affine expansion of $\mathbf{l}_r(\mu)$ or its online evaluation can be too expensive (e.g., when $\mathbf{l}_r(\mu)$ has too many terms in the affine expansion)

and has to be avoided. Therefore, instead of computing the output quantity associated with the solution of the reduced model, we shall approximate it using

$$\mathbf{l}^\Theta(\mu) := \Theta \mathbf{R}_U^{-1} \mathbf{l}(\mu)$$

along with a few additional efficiently computable quantities (see Section 5 for more details). This procedure can also allow an efficient approximation of quadratic quantities of interest and primal-dual corrections.

3 Minimal residual projection

In this section we first present the standard minimal residual projection in a form that allows an easy introduction of random sketching. Then we introduce the sketched version of the minimal residual projection and provide conditions to guarantee its quality.

3.1 Standard minimal residual projection

Let $U_r \subset U$ be a subspace of U (typically obtained with a greedy algorithm or approximate POD). The minres approximation $\mathbf{u}_r(\mu) \in U_r$ of $\mathbf{u}(\mu)$ can be defined by

$$\mathbf{u}_r(\mu) = \arg \min_{\mathbf{w} \in U_r} \|\mathbf{r}(\mathbf{w}; \mu)\|_{U'}. \quad (8)$$

For linear problems it is equivalently characterized by the following (Petrov-)Galerkin orthogonality condition:

$$\langle \mathbf{r}(\mathbf{u}_r(\mu); \mu), \mathbf{w} \rangle = 0, \quad \forall \mathbf{w} \in V_r(\mu), \quad (9)$$

where $V_r(\mu) := \{\mathbf{R}_U^{-1} \mathbf{A}(\mu) \mathbf{x} : \mathbf{x} \in U_r\}$.

If the operator $\mathbf{A}(\mu)$ is invertible then (8) is well-posed. In order to characterize the quality of the projection $\mathbf{u}_r(\mu)$ we define the following parameter-dependent constants

$$\zeta_r(\mu) := \min_{\mathbf{x} \in (\text{span}\{\mathbf{u}(\mu)\} + U_r) \setminus \{0\}} \frac{\|\mathbf{A}(\mu) \mathbf{x}\|_{U'}}{\|\mathbf{x}\|_U}, \quad (10a)$$

$$\iota_r(\mu) := \max_{\mathbf{x} \in (\text{span}\{\mathbf{u}(\mu)\} + U_r) \setminus \{0\}} \frac{\|\mathbf{A}(\mu) \mathbf{x}\|_{U'}}{\|\mathbf{x}\|_U}. \quad (10b)$$

Let $\mathbf{P}_W : U \rightarrow W$ denote the orthogonal projection from U on a subspace $W \subset U$, defined for $\mathbf{x} \in U$ by

$$\mathbf{P}_W \mathbf{x} = \arg \min_{\mathbf{w} \in W} \|\mathbf{x} - \mathbf{w}\|_U.$$

Proposition 3.1. *If $\mathbf{u}_r(\mu)$ satisfies (8) and $\zeta_r(\mu) > 0$, then*

$$\|\mathbf{u}(\mu) - \mathbf{u}_r(\mu)\|_U \leq \frac{\iota_r(\mu)}{\zeta_r(\mu)} \|\mathbf{u}(\mu) - \mathbf{P}_{U_r} \mathbf{u}(\mu)\|_U. \quad (11)$$

Proof. See appendix. □

The constants $\zeta_r(\mu)$ and $\iota_r(\mu)$ can be bounded by the minimal and maximal singular values of $\mathbf{A}(\mu)$:

$$\alpha(\mu) := \min_{\mathbf{x} \in U \setminus \{\mathbf{0}\}} \frac{\|\mathbf{A}(\mu)\mathbf{x}\|_{U'}}{\|\mathbf{x}\|_U} \leq \zeta_r(\mu), \quad (12a)$$

$$\beta(\mu) := \max_{\mathbf{x} \in U \setminus \{\mathbf{0}\}} \frac{\|\mathbf{A}(\mu)\mathbf{x}\|_{U'}}{\|\mathbf{x}\|_U} \geq \iota_r(\mu). \quad (12b)$$

Bounds of $\alpha(\mu)$ and $\beta(\mu)$ can be obtained theoretically [18] or numerically with the successive constraint method [21].

For each μ , the vector $\mathbf{a}_r(\mu) \in \mathbb{K}^r$ such that $\mathbf{u}_r(\mu) = \mathbf{U}_r \mathbf{a}_r(\mu)$ satisfies (9) can be obtained by solving the following reduced (normal) system of equations:

$$\mathbf{A}_r(\mu) \mathbf{a}_r(\mu) = \mathbf{b}_r(\mu), \quad (13)$$

where $\mathbf{A}_r(\mu) = \mathbf{U}_r^H \mathbf{A}(\mu) \mathbf{R}_U^{-1} \mathbf{A}(\mu) \mathbf{U}_r \in \mathbb{K}^{r \times r}$ and $\mathbf{b}_r(\mu) = \mathbf{U}_r^H \mathbf{A}(\mu) \mathbf{R}_U^{-1} \mathbf{b}(\mu) \in \mathbb{K}^r$. The numerical stability of (13) can be ensured through orthogonalization of \mathbf{U}_r similarly as for the classical Galerkin projection. Such orthogonalization yields the following bound for the condition number of $\mathbf{A}_r(\mu)$:

$$\kappa(\mathbf{A}_r(\mu)) := \|\mathbf{A}_r(\mu)\| \|\mathbf{A}_r(\mu)^{-1}\| \leq \left(\frac{\iota_r(\mu)}{\zeta_r(\mu)} \right)^2 \leq \left(\frac{\beta(\mu)}{\alpha(\mu)} \right)^2. \quad (14)$$

This bound can be insufficient for problems with matrix $\mathbf{A}(\mu)$ having a high or even moderate condition number.

The random sketching technique can be used to improve the efficiency and numerical stability of the minimal residual projection, as shown below.

3.2 Sketched minimal residual projection

Let $\Theta \in \mathbb{K}^{k \times n}$ be a certain $U \rightarrow \ell_2$ subspace embedding. The sketched minres projection can be defined by (8) with the dual norm $\|\cdot\|_{U'}$ replaced by its estimation $\|\cdot\|_{U'}^\Theta$, which results in an approximation

$$\mathbf{u}_r(\mu) = \arg \min_{\mathbf{w} \in U_r} \|\mathbf{r}(\mathbf{w}; \mu)\|_{U'}^\Theta. \quad (15)$$

The quasi-optimality of such a projection can be controlled in exactly the same manner as the quasi-optimality of the original minres projection. By defining the constants

$$\zeta_r^\Theta(\mu) := \min_{\mathbf{x} \in (\text{span}\{\mathbf{u}(\mu)\} + U_r) \setminus \{\mathbf{0}\}} \frac{\|\mathbf{A}(\mu)\mathbf{x}\|_{U'}^\Theta}{\|\mathbf{x}\|_U}, \quad (16a)$$

$$\iota_r^\Theta(\mu) := \max_{\mathbf{x} \in (\text{span}\{\mathbf{u}(\mu)\} + U_r) \setminus \{\mathbf{0}\}} \frac{\|\mathbf{A}(\mu)\mathbf{x}\|_{U'}^\Theta}{\|\mathbf{x}\|_U}, \quad (16b)$$

we obtain the following result.

Proposition 3.2. *If $\mathbf{u}_r(\mu)$ satisfies (15) and $\zeta_r^\Theta(\mu) > 0$, then*

$$\|\mathbf{u}(\mu) - \mathbf{u}_r(\mu)\|_U \leq \frac{\iota_r^\Theta(\mu)}{\zeta_r^\Theta(\mu)} \|\mathbf{u}(\mu) - \mathbf{P}_{U_r} \mathbf{u}(\mu)\|_U. \quad (17)$$

Proof. See appendix. □

It follows that if $\zeta_r^\Theta(\mu)$ and $\iota_r^\Theta(\mu)$ are almost equal to $\zeta_r(\mu)$ and $\iota_r(\mu)$, respectively, then the quasi-optimality of the original minres projection (8) shall be almost preserved by its sketched version (15). These properties of $\iota_r^\Theta(\mu)$ and $\zeta_r^\Theta(\mu)$ can be guaranteed under some conditions on Θ (see Proposition 3.3).

Proposition 3.3. *Define the subspace*

$$R_r(U_r; \mu) := \text{span}\{\mathbf{R}_U^{-1} \mathbf{r}(\mathbf{x}; \mu) : \mathbf{x} \in U_r\}. \quad (18)$$

If Θ is a $U \rightarrow \ell_2$ ε -subspace embedding for $R_r(U_r; \mu)$, then

$$\sqrt{1 - \varepsilon} \zeta_r(\mu) \leq \zeta_r^\Theta(\mu) \leq \sqrt{1 + \varepsilon} \zeta_r(\mu), \text{ and } \sqrt{1 - \varepsilon} \iota_r(\mu) \leq \iota_r^\Theta(\mu) \leq \sqrt{1 + \varepsilon} \iota_r(\mu). \quad (19)$$

Proof. See appendix. □

An embedding Θ satisfying an $U \rightarrow \ell_2$ ε -subspace embedding property for the subspace $R_r(U_r; \mu)$ defined in (18), for all $\mu \in \mathcal{P}$ simultaneously, with high probability, may be generated from an oblivious embedding of sufficiently large size. Note that $\dim(R_r(U_r; \mu)) \leq r + 1$. The number of rows k of the oblivious embedding may be selected a priori using the bounds provided in [3], along with a union bound for the probability of success or the fact that $\bigcup_{\mu \in \mathcal{P}} R_r(U_r; \mu)$ is contained in a low-dimensional space. Alternatively, a better value for k can be chosen with a posteriori procedure explained in Section 6. Note that if (19) is satisfied then the quasi-optimality constants of the minres projection are guaranteed to be preserved up to a small factor depending only on the value of ε . Since Θ is here constructed in an oblivious way, the accuracy of random sketching for minres projection can be controlled regardless of the properties of $\mathbf{A}(\mu)$ (e.g., coercivity, condition number, etc.). Recall that in [3] it was revealed that the preservation of the quasi-optimality constants of the classical Galerkin projection by its sketched version is sensitive to the operator's properties. More specifically, random sketching can worsen quasi-optimality constants dramatically for non-coercive or ill-conditioned problems. Therefore, due to its remarkable advantages, the sketched minres projection should be preferred to the sketched Galerkin projection.

Remark 3.4. *Random sketching is not the only way to construct Θ which satisfies the condition in Proposition 3.3 for all $\mu \in \mathcal{P}$. Such a sketching matrix can also be obtained deterministically through approximation of the manifold $R_r^*(U_r) = \{\mathbf{x} \in R_r(U_r, \mu) : \|\mathbf{x}\|_U = 1, \mu \in \mathcal{P}\}$. This approximation can be performed using POD or greedy algorithms. There are two main advantages of random sketching over the deterministic approaches. First, random sketching allows drastic reduction of the computational costs in the offline stage. The second advantage is the oblivious construction of Θ without the knowledge of U_r , which can be particularly important when U_r is*

constructed adaptively (e.g., with a greedy algorithm). Note that the condition in Proposition 3.3 can be satisfied (for not too small ε , say $\varepsilon = 0.1$) with high probability by using an oblivious embedding with $\mathcal{O}(r)$ rows, which is close to the minimal possible value $k = r$. Therefore, the construction of Θ with random sketching in general should be preferred over the deterministic construction.

The vector of coordinates $\mathbf{a}_r(\mu) \in \mathbb{K}^r$ in the basis \mathbf{U}_r of the sketched projection $\mathbf{u}_r(\mu)$ defined by (15) may be obtained in a classical way, i.e., by considering a parameter-dependent reduced (normal) system of equations similar to (13). As for the classical approach, this may lead to numerical instabilities during either the online evaluation of the reduced system from the affine expansions or its solution. A remedy is to directly consider

$$\mathbf{a}_r(\mu) = \arg \min_{\mathbf{x} \in \mathbb{K}^r} \|\mathbf{A}(\mu)\mathbf{U}_r\mathbf{x} - \mathbf{b}(\mu)\|_{U'}^{\Theta} = \arg \min_{\mathbf{x} \in \mathbb{K}^r} \|\mathbf{V}_r^{\Theta}(\mu)\mathbf{x} - \mathbf{b}^{\Theta}(\mu)\|. \quad (20)$$

Since the sketched matrix $\mathbf{V}_r^{\Theta}(\mu)$ and vector $\mathbf{b}^{\Theta}(\mu)$ are of rather small sizes, the minimization problem (20) may be efficiently formed (from the precomputed affine expansions) and then solved (e.g., using QR factorization or SVD) in the online stage.

Proposition 3.5. *If Θ is an ε -embedding for U_r , and \mathbf{U}_r is orthogonal with respect $\langle \cdot, \cdot \rangle_U^{\Theta}$ then the condition number of $\mathbf{V}_r^{\Theta}(\mu)$ is bounded by $\sqrt{\frac{1+\varepsilon}{1-\varepsilon} \frac{t_r^{\Theta}}{\zeta_{\Theta}^{\Theta}}}$.*

Proof. See appendix. □

It follows from Proposition 3.5 (along with Proposition 3.3) that considering (20) can provide better numerical stability than solving reduced systems of equations with standard methods. Furthermore, since affine expansions of $\mathbf{V}_r^{\Theta}(\mu)$ and $\mathbf{b}^{\Theta}(\mu)$ have less terms than affine expansions of $\mathbf{A}_r(\mu)$ and $\mathbf{b}_r(\mu)$ in (13), their online assembling should also be much more stable.

The online efficiency can be further improved with a procedure similar to the one depicted in [3, Section 4.5]. Consider the following oblivious $U \rightarrow \ell_2$ subspace embedding

$$\Phi = \Gamma\Theta,$$

where $\Gamma \in \mathbb{K}^{k' \times k}$, $k' < k$, is a small $(\varepsilon', \delta', r+1)$ oblivious $\ell_2 \rightarrow \ell_2$ subspace embedding. For a given value of the parameter, the solution to (20) can be accurately estimated by

$$\mathbf{u}_r(\mu) \approx \arg \min_{\mathbf{w} \in U_r} \|\mathbf{r}(\mathbf{w}; \mu)\|_{U'}^{\Phi}, \quad (21)$$

with a probability of at least $1 - \delta'$. Note that by [3], $k' = \mathcal{O}(r)$ (in practice, with a small constant, say $k' = 3r$) is enough to provide an accurate estimation of (15) with high probability. For online efficiency, we can use a fixed Θ such that (15) is guaranteed to provide an accurate approximation (see Proposition 3.3) for all $\mu \in \mathcal{P}$ simultaneously, but consider different realizations of a smaller matrix Γ for each particular test set $\mathcal{P}_{\text{test}}$ composed of several parameter values. In this way, in the offline stage a Θ -sketch of \mathbf{U}_r can be precomputed and maintained for the online computations. Thereafter, for the given test set $\mathcal{P}_{\text{test}}$ (with the corresponding new realization

of Γ) the affine expansions of small matrices $\mathbf{V}^\Phi(\mu) := \Gamma\mathbf{V}^\Theta(\mu)$ and $\mathbf{b}^\Phi(\mu) := \Gamma\mathbf{b}^\Theta(\mu)$ can be efficiently precomputed from the Θ -sketch in the “intermediate” online stage. And finally, for each $\mu \in \mathcal{P}_{\text{test}}$, the vector of coordinates of $\mathbf{u}_r(\mu)$ can be obtained by evaluating $\mathbf{V}^\Phi(\mu)$ and $\mathbf{b}^\Phi(\mu)$ from just precomputed affine expansions, and solving

$$\mathbf{a}_r(\mu) = \arg \min_{\mathbf{x} \in \mathbb{K}^r} \|\mathbf{V}_r^\Phi(\mu)\mathbf{x} - \mathbf{b}^\Phi(\mu)\| \quad (22)$$

with a standard method such as QR factorization or SVD.

4 Dictionary-based minimal residual method

Classical RB method becomes ineffective for parameter-dependent problems for which the solution manifold $\mathcal{M} := \{\mathbf{u}(\mu) : \mu \in \mathcal{P}\}$ cannot be well approximated by a single low-dimensional subspace, i.e., its Kolmogorov r -width does not decay rapidly. One can extend the classical RB method by considering a reduced subspace $U_r(\mu)$ depending on a parameter μ . One way to obtain $U_r(\mu)$ is to use a hp -refinement method as in [14, 15], which consists in partitioning the parameter set \mathcal{P} into subsets $\{\mathcal{P}_i\}_{i=1}^M$ and in associating to each subset \mathcal{P}_i a subspace $U_r^i \subset U$ of dimension at most r , therefore resulting in $U_r(\mu) = U_r^i$ if $\mu \in \mathcal{P}_i$, $1 \leq i \leq M$. More formally, the hp -refinement method aims to approximate \mathcal{M} with a library $\mathcal{L}_r := \{U_r^i : 1 \leq i \leq M\}$ of low-dimensional subspaces. For efficiency, the number of subspaces in \mathcal{L}_r has to be moderate (no more than $\mathcal{O}(r^\nu)$ for some small ν , say $\nu = 2$ or 3 , which should be dictated by the particular computational architecture). A nonlinear Kolmogorov r -width of \mathcal{M} with a library of M subspaces can be defined as in [31] by

$$d_r(\mathcal{M}; M) = \inf_{\#\mathcal{L}_r=M} \sup_{\mathbf{u} \in \mathcal{M}} \min_{W_r \in \mathcal{L}_r} \|\mathbf{u} - \mathbf{P}_{W_r}\mathbf{u}\|_U, \quad (23)$$

where the infimum is taken over all libraries of M subspaces. Clearly, the approximation $\mathbf{P}_{U_r(\mu)}\mathbf{u}(\mu)$ over a parameter-dependent subspace $U_r(\mu)$ associated with a partitioning of \mathcal{P} into M subdomains satisfies

$$d_r(\mathcal{M}; M) \leq \max_{\mu \in \mathcal{P}} \|\mathbf{u}(\mu) - \mathbf{P}_{U_r(\mu)}\mathbf{u}(\mu)\|_U. \quad (24)$$

Therefore, for the hp -refinement method to be effective, the solution manifold is required to be well approximable in terms of the measure $d_r(\mathcal{M}; M)$.

The hp -refinement method may present serious drawbacks: it can be highly sensitive to the parametrization, it can require a large number of subdomains in \mathcal{P} (especially for high-dimensional parameter domains) and it can require computing too many solution samples. These drawbacks can be partially reduced by various modifications of the hp -refinement method [2, 26], but not circumvented.

We here propose a dictionary-based method, which can be seen as an alternative to a partitioning of \mathcal{P} for defining $U_r(\mu)$, and argue why this method is more natural and can be applied to a larger class of problems.

4.1 Dictionary-based approximation

For each value μ of the parameter, the basis vectors for $U_r(\mu)$ are selected online from a certain dictionary \mathcal{D}_K of K candidate vectors in U , $K \geq r$. For efficiency of the algorithms in the particular computational environment, the value for K has to be chosen as $\mathcal{O}(r^\nu)$ with a small ν similarly as the number of subdomains M for the hp -refinement method. Let $\mathcal{L}_r(\mathcal{D}_K)$ denote the library of all subspaces spanned by r vectors from \mathcal{D}_K . A dictionary-based r -width is defined as

$$\sigma_r(\mathcal{M}; K) = \inf_{\#\mathcal{D}_K=K} \sup_{\mathbf{u} \in \mathcal{M}} \min_{W_r \in \mathcal{L}_r(\mathcal{D}_K)} \|\mathbf{u} - \mathbf{P}_{W_r} \mathbf{u}\|_U, \quad (25)$$

where the infimum is taken over all subsets \mathcal{D}_K of U with cardinality $\#\mathcal{D}_K = K$. A dictionary \mathcal{D}_K can be efficiently constructed offline with an adaptive greedy procedure (see Section 4.4).

In general, the performance of the method can be characterized through the approximability of the solution manifold \mathcal{M} in terms of the r -width, and quasi-optimality of the considered $U_r(\mu)$ compared to the best approximation. The dictionary-based approximation can be beneficial over the refinement methods in either of these aspects, which is explained below.

It can be easily shown that

$$\sigma_r(\mathcal{M}; K) \leq d_r(\mathcal{M}; M), \text{ for } K \geq rM.$$

Therefore, if a solution manifold can be well approximated with a partitioning of the parameter domain into M subdomains each associated with a subspace of dimension r , then it should also be well approximated with a dictionary of size $K = rM$, which implies a similar computational cost. The converse statement, however, is not true. A dictionary with K vectors can generate a library with up to $\binom{K}{r}$ subspaces so that

$$d_r(\mathcal{M}; \binom{K}{r}) \leq \sigma_r(\mathcal{M}; K).$$

Consequently, to obtain a decay of $d_r(\mathcal{M}; M)$ with r similar to the decay of $\sigma_r(\mathcal{M}; r^\nu)$, we can be required to use M which depends exponentially on r .

The great potential of the dictionary-based approximation can be justified by important properties of the dictionary-based r -width given in Proposition 4.1 and Corollary 4.2.

Proposition 4.1. *Let \mathcal{M} be obtained by the superposition of parameter-dependent vectors:*

$$\mathcal{M} = \left\{ \sum_{i=1}^l \mathbf{u}^{(i)}(\mu) : \mu \in \mathcal{P} \right\}, \quad (26)$$

where $\mathbf{u}^{(i)}(\mu) \in U$, $i = 1, \dots, l$. Then, we have

$$\sigma_r(\mathcal{M}; K) \leq \sum_{i=1}^l \sigma_{r_i}(\mathcal{M}^{(i)}; K_i), \quad (27)$$

with $r = \sum_{i=1}^l r_i$, $K = \sum_{i=1}^l K_i$ and

$$\mathcal{M}^{(i)} = \{ \mathbf{u}^{(i)}(\mu) : \mu \in \mathcal{P} \}. \quad (28)$$

Proof. See appendix. □

Corollary 4.2 (Approximability of a superposition of solutions). *Consider several solution manifolds $\mathcal{M}^{(i)}$ defined by (28), $1 \leq i \leq l$, and the resulting manifold \mathcal{M} defined by (26). Let c , C , α , β and γ be some constants. The following properties hold.*

- (i) *If $\sigma_r(\mathcal{M}^{(i)}; cr^\nu) \leq Cr^{-\alpha}$, then $\sigma_r(\mathcal{M}; cl^{1-\nu}r^\nu) \leq Cl^{1+\alpha}r^{-\alpha}$,*
- (ii) *If $\sigma_r(\mathcal{M}^{(i)}; cr^\nu) \leq Ce^{-\gamma r^\beta}$, then $\sigma_r(\mathcal{M}; cl^{1-\nu}r^\nu) \leq Cle^{-\gamma l^{-\beta}r^\beta}$.*

From Proposition 4.1 and Corollary 4.2 it follows that the approximability of the solution manifold in terms of the dictionary-based r -width is preserved under the superposition operation. In other words, if the dictionary-based r -widths of manifolds $\mathcal{M}^{(i)}$ have a certain decay with r (e.g., exponential or algebraic), by using dictionaries containing $K = \mathcal{O}(r^\nu)$ vectors, then the type of decay is preserved by their superposition (with the same rate for the algebraic decay). This property can be crucial for problems where the solution is a superposition of several contributions (possibly unknown), which is a quite typical situation. For instance, we have such a situation for PDEs with multiple transport phenomena. A similar property as (27) also holds for the classical linear Kolmogorov r -width $d_r(\mathcal{M})$. Namely, we have

$$d_r(\mathcal{M}) \leq \sum_{i=1}^l d_{r_i}(\mathcal{M}^{(i)}), \quad (29)$$

with $r = \sum_{i=1}^l r_i$. This relation follows immediately from Proposition 4.1 and the fact that $d_r(\mathcal{M}) = \sigma_r(\mathcal{M}; r)$. For the nonlinear Kolmogorov r -width (24), however, the relation

$$d_r(\mathcal{M}, M) \leq \sum_{i=1}^l d_{r_i}(\mathcal{M}^{(i)}, M^{(i)}), \quad (30)$$

where $r = \sum_{i=1}^l r_i$, holds under the condition that $M \geq \prod_{i=1}^l M^{(i)}$. In general, the preservation of the type of decay with r of $d_r(\mathcal{M}, M)$, by using libraries with $M = \mathcal{O}(r^\nu)$ terms, may not be guaranteed. It can require libraries with much larger numbers of r -dimensional subspaces than $\mathcal{O}(r^\nu)$, namely $M = \mathcal{O}(r^{l\nu})$ subspaces.

Another advantage of the dictionary-based method is its weak sensitivity to the parametrization of the manifold \mathcal{M} , in contrast to the hp -refinement method, for which a bad choice of parametrization can result in approximations with too many local reduced subspaces. Indeed, the solution map $\mu \rightarrow \mathbf{u}(\mu)$ is often expected to have certain properties (e.g., symmetries or anisotropies) that yield the existence of a better parametrization of \mathcal{M} than the one proposed by the user. Finding a good parametrization of the solution manifold may require a deep intrusive analysis of the problem, and is therefore usually an intractable task. On the other hand, our dictionary-based methodology provides a reduced approximation subspace for each vector from \mathcal{M} regardless of the chosen parametrization.

4.2 Sparse minimal residual approximation

Here we assume to be given a dictionary \mathcal{D}_K of K vectors in U . Ideally, for each μ , $\mathbf{u}(\mu)$ should be approximated by orthogonal projection onto a subspace $W_r(\mu)$ that minimizes

$$\|\mathbf{u}(\mu) - \mathbf{P}_{W_r(\mu)}\mathbf{u}(\mu)\|_U \quad (31)$$

over the library $\mathcal{L}_r(\mathcal{D}_K)$. The selection of the optimal subspace requires operating with the exact solution $\mathbf{u}(\mu)$ which is prohibited. Therefore, the reduced approximation space $U_r(\mu)$ and the associated approximate solution $\mathbf{u}_r(\mu) \in U_r(\mu)$ are defined such that

$$U_r(\mu) \in \arg \min_{W_r \in \mathcal{L}_r(\mathcal{D}_K)} \min_{\mathbf{w} \in W_r} \|\mathbf{r}(\mathbf{w}; \mu)\|_{U'}, \quad \mathbf{u}_r(\mu) = \arg \min_{\mathbf{w} \in U_r(\mu)} \|\mathbf{r}(\mathbf{w}; \mu)\|_{U'}. \quad (32)$$

The solution $\mathbf{u}_r(\mu)$ from (32) shall be referred to as sparse minres approximation (relatively to the dictionary \mathcal{D}_K). The quasi-optimality of this approximation can be characterized with the following parameter-dependent constants:

$$\zeta_{r,K}(\mu) := \min_{W_r \in \mathcal{L}_r(\mathcal{D}_K)} \min_{\mathbf{x} \in (\text{span}\{\mathbf{u}(\mu)\} + W_r) \setminus \{\mathbf{0}\}} \frac{\|\mathbf{A}(\mu)\mathbf{x}\|_{U'}}{\|\mathbf{x}\|_U}, \quad (33a)$$

$$\iota_{r,K}(\mu) := \max_{W_r \in \mathcal{L}_r(\mathcal{D}_K)} \max_{\mathbf{x} \in (\text{span}\{\mathbf{u}(\mu)\} + W_r) \setminus \{\mathbf{0}\}} \frac{\|\mathbf{A}(\mu)\mathbf{x}\|_{U'}}{\|\mathbf{x}\|_U}. \quad (33b)$$

In general, one can bound $\zeta_{r,K}(\mu)$ and $\iota_{r,K}(\mu)$ by the minimal and the maximal singular values $\alpha(\mu)$ and $\beta(\mu)$ of $\mathbf{A}(\mu)$. Observe also that for $K = r$ (i.e., when the library $\mathcal{L}_r(\mathcal{D}_K) = \{U_r\}$ has a single subspace) we have $\zeta_{r,K}(\mu) = \zeta_r(\mu)$ and $\iota_{r,K}(\mu) = \iota_r(\mu)$.

Proposition 4.3. *Let $\mathbf{u}_r(\mu)$ be the solution of (32) and $\zeta_{r,K}(\mu) > 0$, then*

$$\|\mathbf{u}(\mu) - \mathbf{u}_r(\mu)\|_U \leq \frac{\iota_{r,K}(\mu)}{\zeta_{r,K}(\mu)} \min_{W_r \in \mathcal{L}_r(\mathcal{D}_K)} \|\mathbf{u}(\mu) - \mathbf{P}_{W_r}\mathbf{u}(\mu)\|_U. \quad (34)$$

Proof. See appendix. □

Let $\mathbf{U}_K \in \mathbb{K}^{n \times K}$ be a matrix whose columns are the vectors in the dictionary \mathcal{D}_K and $\mathbf{a}_{r,K}(\mu) \in \mathbb{K}^K$, with $\|\mathbf{a}_{r,K}(\mu)\|_0 \leq r$, be the r -sparse vector of coordinates of $\mathbf{u}_r(\mu)$ in the dictionary, i.e., $\mathbf{u}_r(\mu) = \mathbf{U}_K \mathbf{a}_{r,K}(\mu)$. The vector of coordinates associated with the solution $\mathbf{u}_r(\mu)$ of (32) is the solution to the following parameter-dependent sparse least-squares problem:

$$\min_{\mathbf{z} \in \mathbb{K}^K} \|\mathbf{A}(\mu)\mathbf{U}_K\mathbf{z} - \mathbf{b}(\mu)\|_{U'}, \quad \text{subject to } \|\mathbf{z}\|_0 \leq r. \quad (35)$$

For each $\mu \in \mathcal{P}$ an approximate solution to problem (35) can be obtained with a standard greedy algorithm depicted in Algorithm 1. It selects the nonzero entries of $\mathbf{a}_{r,K}(\mu)$ one by one to minimize the residual. The algorithm corresponds to either the orthogonal greedy (also called Orthogonal Matching Pursuit in signal processing community [32]) or stepwise projection algorithm (see [12]) depending on whether the (optional) Step 8 (which is the orthogonalization of $\{\mathbf{v}_j(\mu)\}_{j=1}^K$ with respect to

$V_i(\mu)$) is considered. It should be noted that performing Step 8 can be of great importance due to possible high mutual coherence of the (mapped) dictionary $\{\mathbf{v}_j(\mu)\}_{j=1}^K$. Algorithm 1 is provided in a conceptual form. A more sophisticated procedure can be derived to improve the online efficiency (e.g., considering precomputed affine expansions of $\mathbf{A}_K(\mu) := \mathbf{U}_K^H \mathbf{A}(\mu)^H \mathbf{R}_U^{-1} \mathbf{A}(\mu) \mathbf{U}_K \in \mathbb{K}^{K \times K}$ and $\mathbf{b}_K(\mu) = \mathbf{U}_K^H \mathbf{A}(\mu)^H \mathbf{R}_U^{-1} \mathbf{b}(\mu) \in \mathbb{K}^K$, updating the residual using a Gram-Schmidt procedure, etc). Algorithm 1, even when efficiently implemented, can still require heavy computations in both the offline and online stages, and be numerically unstable. One of the contributions of this paper is a drastic improvement of its efficiency and stability by random sketching, thus making the use of dictionary-based model reduction feasible in practice.

Algorithm 1 Orthogonal greedy algorithm

Given: μ , $\mathbf{U}_K = [\mathbf{w}_j]_{j=1}^K$, $\mathbf{A}(\mu)$, $\mathbf{b}(\mu)$, τ , r .

Output: index set $\Lambda_r(\mu)$, the coordinates $\mathbf{a}_r(\mu)$ of $\mathbf{u}_r(\mu)$ on the basis $\{\mathbf{w}_j\}_{j \in \Lambda_r(\mu)}$.

1. Set $i := 0$, $U_0(\mu) = \{\mathbf{0}\}$, $\mathbf{u}_0(\mu) = \mathbf{0}$, $\Lambda_0(\mu) = \emptyset$, $\tilde{\Delta}_0(\mu) = \infty$.

2. Set $[\mathbf{v}_1(\mu), \dots, \mathbf{v}_K(\mu)] := \mathbf{A}(\mu) \mathbf{U}_K$ and normalize the vectors $\mathbf{v}_j(\mu)$, $1 \leq j \leq K$.

while $\tilde{\Delta}_i(\mu) \geq \tau$ and $i < r$ **do**

3. Set $i := i + 1$.

4. Find the index $p_i \in \{1, \dots, K\}$ which maximizes $|\langle \mathbf{v}_{p_i}(\mu), \mathbf{r}(\mathbf{u}_{i-1}(\mu); \mu) \rangle_{U'}|$.

5. Set $\Lambda_i(\mu) := \Lambda_{i-1}(\mu) \cup \{p_i\}$.

6. Solve (13) with a reduced matrix $\mathbf{U}_i(\mu) = [\mathbf{w}_j]_{j \in \Lambda_i(\mu)}$ and obtain the coordinates $\mathbf{a}_i(\mu)$.

7. Compute error bound $\tilde{\Delta}_i(\mu)$ of $\mathbf{u}_i(\mu) = \mathbf{U}_i(\mu) \mathbf{a}_i(\mu)$.

8. (Optional) Set $\mathbf{v}_j(\mu) := \mathbf{v}_j(\mu) - \mathbf{P}_{V_i(\mu)} \mathbf{v}_j(\mu)$, where $\mathbf{P}_{V_i(\mu)}$ is the orthogonal projector on $V_i(\mu) := \text{span}(\{\mathbf{v}_p(\mu)\}_{p \in \Lambda_i(\mu)})$, and normalize $\mathbf{v}_j(\mu)$, $1 \leq j \leq K$.

end while

4.3 Sketched sparse minimal residual approximation

Let $\Theta \in \mathbb{K}^{k \times n}$ be a certain $U \rightarrow \ell_2$ subspace embedding. The sparse minres approximation defined by (32), associated with dictionary \mathcal{D}_K , can be estimated by the solution $\mathbf{u}_r(\mu)$ of the following minimization problem

$$U_r(\mu) \in \arg \min_{W_r \in \mathcal{L}_r(\mathcal{D}_K)} \min_{\mathbf{w} \in W_r} \|\mathbf{r}(\mathbf{w}; \mu)\|_{U'}^\Theta, \quad \mathbf{u}_r(\mu) = \arg \min_{\mathbf{w} \in U_r(\mu)} \|\mathbf{r}(\mathbf{w}; \mu)\|_{U'}^\Theta. \quad (36)$$

In order to characterize the quasi-optimality of the sketched sparse minres approximation defined by (36) we introduce the following parameter-dependent values

$$\zeta_{r,K}^\Theta(\mu) := \min_{W_r \in \mathcal{L}_r(\mathcal{D}_K)} \min_{\mathbf{x} \in (\text{span}\{\mathbf{u}(\mu)\} + W_r) \setminus \{\mathbf{0}\}} \frac{\|\mathbf{A}(\mu) \mathbf{x}\|_{U'}^\Theta}{\|\mathbf{x}\|_U}, \quad (37a)$$

$$\iota_{r,K}^\Theta(\mu) := \max_{W_r \in \mathcal{L}_r(\mathcal{D}_K)} \max_{\mathbf{x} \in (\text{span}\{\mathbf{u}(\mu)\} + W_r) \setminus \{\mathbf{0}\}} \frac{\|\mathbf{A}(\mu) \mathbf{x}\|_{U'}^\Theta}{\|\mathbf{x}\|_U}. \quad (37b)$$

Observe that choosing $K = r$ yields $\zeta_{r,K}^\Theta(\mu) = \zeta_r^\Theta(\mu)$ and $\iota_{r,K}^\Theta(\mu) = \iota_r^\Theta(\mu)$.

Proposition 4.4. *If $\mathbf{u}_r(\mu)$ satisfies (36) and $\zeta_{r,K}^\Theta(\mu) > 0$, then*

$$\|\mathbf{u}(\mu) - \mathbf{u}_r(\mu)\|_U \leq \frac{\iota_{r,K}^\Theta(\mu)}{\zeta_{r,K}^\Theta(\mu)} \min_{W_r \in \mathcal{L}_r(\mathcal{D}_K)} \|\mathbf{u}(\mu) - \mathbf{P}_{W_r} \mathbf{u}(\mu)\|_U, \quad (38)$$

Proof. See appendix. □

It follows from Proposition 4.4 that the quasi-optimality of the sketched sparse minres approximation can be controlled by bounding the constants $\zeta_{r,K}^\Theta(\mu)$ and $\iota_{r,K}^\Theta(\mu)$.

Proposition 4.5. *If Θ is a $U \rightarrow \ell_2$ ε -embedding for every subspace $R_r(W_r; \mu)$, defined by (18), with $W_r \in \mathcal{L}_r(\mathcal{D}_K)$, then*

$$\sqrt{1 - \varepsilon} \zeta_{r,K}(\mu) \leq \zeta_{r,K}^\Theta(\mu) \leq \sqrt{1 + \varepsilon} \zeta_{r,K}(\mu), \quad (39a)$$

and

$$\sqrt{1 - \varepsilon} \iota_{r,K}(\mu) \leq \iota_{r,K}^\Theta(\mu) \leq \sqrt{1 + \varepsilon} \iota_{r,K}(\mu). \quad (39b)$$

Proof. See appendix. □

By Definition 2.3 and the union bound for the probability of success, if Θ is a $(\varepsilon, \binom{K}{r}^{-1} \delta, r+1)$ oblivious $U \rightarrow \ell_2$ subspace embedding, then Θ satisfies the assumption of Proposition 4.5 with probability of at least $1 - \delta$. The sufficient number of rows for Θ may be chosen a priori with the bounds provided in [3] or adaptively with a procedure from Section 6. For the Gaussian embeddings the a priori bounds are logarithmic in K and n , and proportional to r . For SRHT they are also logarithmic in K and n , but proportional to r^2 (although in practice SRHT performs equally well as the Gaussian distribution). Moreover, if \mathcal{P} is a finite set, an oblivious embedding Θ which satisfies the hypothesis of Proposition 4.5 for all $\mu \in \mathcal{P}$, simultaneously, may be chosen using the above considerations and a union bound for the probability of success. Alternatively, for an infinite set \mathcal{P} , Θ can be chosen as an ε -embedding for a collection of low-dimensional subspaces $R_r^*(W_r)$ (which can be obtained from the affine expansions of $\mathbf{A}(\mu)$ and $\mathbf{b}(\mu)$) each containing $\bigcup_{\mu \in \mathcal{P}} R_r(\mu; W_r)$ and associated with a subspace W_r of $\mathcal{L}_r(\mathcal{D}_K)$. Such an embedding can be again generated in an oblivious way by considering Definition 2.3 and a union bound for the probability of success.

From an algebraic point of view, the optimization problem (36) can be formulated as the following sparse least-squares problem:

$$\min_{\substack{\mathbf{z} \in \mathbb{K}^K \\ \|\mathbf{z}\|_0 \leq r}} \|\mathbf{A}(\mu) \mathbf{U}_K \mathbf{z} - \mathbf{b}(\mu)\|_{U'}^\Theta = \min_{\substack{\mathbf{z} \in \mathbb{K}^K \\ \|\mathbf{z}\|_0 \leq r}} \|\mathbf{V}_K^\Theta(\mu) \mathbf{z} - \mathbf{b}^\Theta(\mu)\|, \quad (40)$$

where $\mathbf{V}_K^\Theta(\mu)$ and $\mathbf{b}^\Theta(\mu)$ are the components (7) of the Θ -sketch of \mathbf{U}_K (a matrix whose columns are the vectors in \mathcal{D}_K). The solution $\mathbf{a}_{r,K}(\mu)$ of (40) is the r -sparse vector of the coordinates of $\mathbf{u}_r(\mu)$. We observe that (40) is simply an approximation of a small vector $\mathbf{b}^\Theta(\mu)$ with a dictionary composed from column vectors of $\mathbf{V}_K^\Theta(\mu)$. Therefore, unlike the original sparse least-squares

problem (35), the solution to its sketched version (40) can be efficiently approximated with standard tools in the online stage. For instance, we can use Algorithm 1 replacing $\langle \cdot, \cdot \rangle_{U'}$ with $\langle \cdot, \cdot \rangle_{U'}^\Theta$. Clearly, in Algorithm 1 the inner products $\langle \cdot, \cdot \rangle_{U'}^\Theta$ should be efficiently evaluated from $\mathbf{V}_K^\Theta(\mu)$ and $\mathbf{b}^\Theta(\mu)$. For this a Θ -sketch of \mathbf{U}_K can be precomputed in the offline stage and then used for online evaluation of $\mathbf{V}_K^\Theta(\mu)$ and $\mathbf{b}^\Theta(\mu)$ for each value of the parameter. Another way to obtain an approximate solution to the sketched sparse least-squares problem (40) is to use LASSO or similar methods.

Let us now characterize the algebraic stability (i.e., sensitivity to round-off errors) of the (approximate) solution of (40). The solution of (40) is essentially obtained from the following least-squares problem

$$\min_{\mathbf{x} \in \mathbb{K}^r} \|\mathbf{V}_r^\Theta(\mu)\mathbf{x} - \mathbf{b}^\Theta(\mu)\|, \quad (41)$$

where $\mathbf{V}_r^\Theta(\mu)$ is a matrix whose column vectors are (adaptively) selected from the columns of $\mathbf{V}_K^\Theta(\mu)$. The algebraic stability of this problem can be measured by the condition number of $\mathbf{V}_r^\Theta(\mu)$. The minimal and the maximal singular values of $\mathbf{V}_r^\Theta(\mu)$ can be bounded using the parameter-dependent coefficients $\iota_{r,K}^\Theta(\mu)$, $\zeta_{r,K}^\Theta(\mu)$ and the so-called restricted isometry property (RIP) constants associated with the dictionary \mathcal{D}_K , which are defined by

$$\Sigma_{r,K}^{\min} := \min_{\substack{\mathbf{z} \in \mathbb{K}^K \\ \|\mathbf{z}\|_0 \leq r}} \frac{\|\mathbf{U}_K \mathbf{z}\|_U}{\|\mathbf{z}\|}, \quad \Sigma_{r,K}^{\max} := \max_{\substack{\mathbf{z} \in \mathbb{K}^K \\ \|\mathbf{z}\|_0 \leq r}} \frac{\|\mathbf{U}_K \mathbf{z}\|_U}{\|\mathbf{z}\|}. \quad (42)$$

Proposition 4.6. *The minimal singular value of $\mathbf{V}_r^\Theta(\mu)$ in (41) is bounded below by $\zeta_{r,K}^\Theta(\mu)\Sigma_{r,K}^{\min}$, while the maximal singular value of $\mathbf{V}_r^\Theta(\mu)$ is bounded above by $\iota_{r,K}^\Theta(\mu)\Sigma_{r,K}^{\max}$.*

Proof. See appendix. □

The RIP constants quantify the linear dependency of the dictionary vectors. For instance, it is easy to see that for a dictionary composed of orthogonal unit vectors we have $\Sigma_{r,K}^{\min} = \Sigma_{r,K}^{\max} = 1$. From Proposition 4.6, one can deduce the maximal level of degeneracy of \mathcal{D}_K for which the sparse optimization problem (40) remains sufficiently stable.

Remark 4.7. *In general, our approach is more stable than the algorithms from [13, 23]. These algorithms basically proceed with the solution of the reduced system of equations $\mathbf{A}_r(\mu)\mathbf{a}_r(\mu) = \mathbf{b}_r(\mu)$, where $\mathbf{A}_r(\mu) = \mathbf{U}_r(\mu)^H \mathbf{A}(\mu) \mathbf{U}_r(\mu)$, with $\mathbf{U}_r(\mu)$ being a matrix whose column vectors are selected from the column vectors of \mathbf{U}_K . In this case, the bounds for the minimal and the maximal singular values of $\mathbf{A}_r(\mu)$ are proportional to the squares of the minimal and the maximal singular values of $\mathbf{U}_r(\mu)$, which implies a quadratic dependency on the RIP constants $\Sigma_{r,K}^{\min}$ and $\Sigma_{r,K}^{\max}$. On the other hand, with (sketched) minres methods the dependency of the singular values of the reduced matrix $\mathbf{V}^\Theta(\mu)$ on $\Sigma_{r,K}^{\min}$ and $\Sigma_{r,K}^{\max}$ is only linear (see Proposition 4.6). Consequently, our methodology provides an improvement of not only efficiency but also numerical stability for problems with high linear dependency of dictionary vectors.*

Similarly to the sketched minres projection, a better online efficiency can be obtained by introducing

$$\Phi = \Gamma \Theta,$$

where $\mathbf{\Gamma} \in \mathbb{K}^{k' \times k}$, $k' < k$, is a small $(\varepsilon', \binom{K}{r}^{-1} \delta', r+1)$ oblivious $\ell_2 \rightarrow \ell_2$ subspace embedding, and approximating the solution to (36) by

$$U_r(\mu) \in \arg \min_{W_r \in \mathcal{L}_r(\mathcal{D}_K)} \min_{\mathbf{w} \in W_r} \|\mathbf{r}(\mathbf{w}; \mu)\|_{U'}^{\Phi}, \quad \mathbf{u}_r(\mu) = \arg \min_{\mathbf{w} \in U_r(\mu)} \|\mathbf{r}(\mathbf{w}; \mu)\|_{U'}^{\Phi}. \quad (43)$$

It follows that the accuracy (and the stability) of the solution of (43) is almost the same as the one of (36) with probability at least $1 - \delta'$. In an algebraic setting, (43) can be expressed as

$$\min_{\substack{\mathbf{z} \in \mathbb{K}^K \\ \|\mathbf{z}\|_0 \leq r}} \|\mathbf{V}_K^{\Phi}(\mu) \mathbf{z} - \mathbf{b}^{\Phi}(\mu)\|, \quad (44)$$

whose solution $\mathbf{a}_r(\mu)$ is a r -sparse vector of coordinates of $\mathbf{u}_r(\mu)$. An approximate solution to such a problem can be computed with Algorithm 1 by replacing $\langle \cdot, \cdot \rangle_{U'}$ with $\langle \cdot, \cdot \rangle_{U'}^{\Phi}$. An efficient procedure for evaluating the coordinates of a sketched dictionary-based approximation on a test set $\mathcal{P}_{\text{test}}$ from the Θ -sketch of \mathbf{U}_K is provided in Algorithm 2. Algorithm 2 uses a residual-based sketched error estimator from [3] defined by

$$\tilde{\Delta}_i(\mu) = \Delta^{\Phi}(\mathbf{u}_i(\mu); \mu) = \frac{\|\mathbf{r}(\mathbf{u}_i(\mu); \mu)\|_{U'}^{\Phi}}{\eta(\mu)}, \quad (45)$$

where $\eta(\mu)$ is a computable lower bound of the minimal singular value of $\mathbf{A}(\mu)$. Let us underline the importance of performing Step 8 (orthogonalization of the dictionary vectors with respect to the previously selected basis vectors), for problems with “degenerate” dictionaries (with high mutual coherence). It should be noted that at Steps 7 and 8 we use a Gram-Schmidt procedure for orthogonalization because of its simplicity and efficiency, whereas a modified Gram-Schmidt algorithm could provide better accuracy. It is also important to note that Algorithm 2 satisfies a basic consistency property in the sense that it exactly recovers the vectors from the dictionary with high probability.

If \mathcal{P} is a finite set, then the theoretical bounds for Gaussian matrices and the empirical experience for SRHT state that choosing $k = \mathcal{O}(r \log K + \log \delta + \log(\#\mathcal{P}))$ and $k' = \mathcal{O}(r \log K + \log \delta + \log(\#\mathcal{P}_{\text{test}}))$ in Algorithm 2 yield a quasi-optimal solution to (31) for all $\mu \in \mathcal{P}_{\text{test}}$ with probability at least $1 - \delta$. Let us neglect the logarithmic summands. Assuming $\mathbf{A}(\mu)$ and $\mathbf{b}(\mu)$ admit affine representations with m_A and m_b terms, it follows that the online complexity and memory consumption of Algorithm 2 is only $\mathcal{O}((m_A K + m_b) r \log K + r^2 K \log K) \#\mathcal{P}_{\text{test}}$ and $\mathcal{O}((m_A K + m_b) r \log K)$, respectively. The quasi-optimality for infinite \mathcal{P} can be ensured with high probability by increasing k to $\mathcal{O}(r^* \log K + \log \delta)$, where r^* is the maximal dimension of subspaces $R_r^*(W_r)$ containing $\bigcup_{\mu \in \mathcal{P}} R_r(\mu; W_r)$ with $W_r \in \mathcal{L}_r(\mathcal{D}_K)$. This shall increase the memory consumption by a factor of r^*/r but should have a negligible effect (especially for large $\mathcal{P}_{\text{test}}$) on the complexity, which is mainly characterized by the size of Φ . Note that for parameter-separable problems we have $r^* \leq m_A r + m_b$.

Algorithm 2 Efficient/stable sketched orthogonal greedy algorithm

Given: $\mathcal{P}_{\text{test}}$, Θ -sketch of $\mathbf{U}_K = [\mathbf{w}_j]_{j=1}^K$, τ , r .

Output: index set $\Lambda_r(\mu)$, the coordinates $\mathbf{a}_r(\mu)$ of $\mathbf{u}_r(\mu)$ on basis $\{\mathbf{w}_j\}_{j \in \Lambda_r(\mu)}$,
and the error indicator $\Delta^\Phi(\mathbf{u}_r(\mu); \mu)$ for each $\mu \in \mathcal{P}_{\text{test}}$.

1. Generate $\mathbf{\Gamma}$ and evaluate the affine factors of $\mathbf{V}_K^\Phi(\mu) := \mathbf{\Gamma} \mathbf{V}_K^\Theta(\mu)$ and $\mathbf{b}^\Phi(\mu) := \mathbf{\Gamma} \mathbf{b}^\Theta(\mu)$.

for $\mu \in \mathcal{P}_{\text{test}}$ **do**

2. Evaluate $[\mathbf{v}_1^\Phi(\mu), \dots, \mathbf{v}_K^\Phi(\mu)] := \mathbf{V}_K^\Phi(\mu)$ and $\mathbf{b}^\Phi(\mu)$ from the affine expansions
and normalize $\mathbf{v}_j^\Phi(\mu)$, $1 \leq j \leq K$.

3. Set $i = 0$, obtain $\eta(\mu)$ in (45), set $\Lambda_0(\mu) = \emptyset$, $\mathbf{r}_0^\Phi(\mu) := \mathbf{b}^\Phi(\mu)$ and $\Delta^\Phi(\mu) := \|\mathbf{b}^\Phi(\mu)\|/\eta(\mu)$.

while $\Delta^\Phi(\mathbf{u}_i(\mu); \mu) \geq \tau$ and $i \leq r$ **do**

4. Set $i := i + 1$.

5. Find the index p_i which maximizes $|\mathbf{v}_{p_i}^\Phi(\mu)^H \mathbf{r}_{i-1}^\Phi(\mu)|$. Set $\Lambda_i(\mu) := \Lambda_{i-1}(\mu) \cup \{p_i\}$.

6. Set $\mathbf{v}_{p_i}^\Phi(\mu) := \mathbf{v}_{p_i}^\Phi(\mu) - \sum_{j=1}^{i-1} \mathbf{v}_{p_j}^\Phi(\mu) [\mathbf{v}_{p_j}^\Phi(\mu)^H \mathbf{v}_{p_i}^\Phi(\mu)]$ and normalize it.

7. Compute $\mathbf{r}_i^\Phi(\mu) := \mathbf{r}_{i-1}^\Phi(\mu) - \mathbf{v}_{p_i}^\Phi(\mu) [\mathbf{v}_{p_i}^\Phi(\mu)^H \mathbf{r}_{i-1}^\Phi(\mu)]$ and $\Delta^\Phi(\mathbf{u}_i(\mu); \mu) = \|\mathbf{r}_i^\Phi(\mu)\|/\eta(\mu)$.

8. (Optional) Set $\mathbf{v}_j^\Phi(\mu) = \mathbf{v}_j^\Phi(\mu) - \mathbf{v}_{p_i}^\Phi(\mu) [\mathbf{v}_{p_i}^\Phi(\mu)^H \mathbf{v}_j^\Phi(\mu)]$ and normalize it, $1 \leq j \leq K$.

end while

9. Solve (22) choosing $r := i$, and the columns p_1, p_2, \dots, p_i of $\mathbf{V}_K^\Phi(\mu)$ as the columns
for $\mathbf{V}_r^\Phi(\mu)$ and obtain solution $\mathbf{a}_r(\mu)$.

end for

4.4 Dictionary generation

The simplest way is to choose the dictionary as a set of solution samples (snapshots) associated with a training set $\mathcal{P}_{\text{train}}$, i.e.,

$$\mathcal{D}_K = \{\mathbf{u}(\mu) : \mu \in \mathcal{P}_{\text{train}}\}. \quad (46)$$

Let us recall that we are interested in computing a Θ -sketch of \mathbf{U}_K (matrix whose columns form \mathcal{D}_K) rather than the full matrix. In certain computational environments, a Θ -sketch of \mathbf{U}_K can be computed very efficiently. For instance, each snapshot may be computed and sketched on a separate distributed machine. Thereafter small sketches can be efficiently transferred to the master workstation for constructing the reduced order model.

A better dictionary may be computed with the greedy procedure presented in Algorithm 3, recursively enriching the dictionary with a snapshot at the parameter value associated with the maximal error at the previous iteration. The value for r (the dimension of the parameter-dependent reduced subspace $U_r(\mu)$) should be chosen according to the particular computational architecture. Since the provisional online solver (identified with Algorithm 2) guarantees exact recovery of snapshots belonging to \mathcal{D}_i , Algorithm 3 is consistent. It has to be noted that the first r iterations of the proposed greedy algorithm for the dictionary generation coincide with the first r iterations of the standard greedy algorithm for the reduced basis generation.

By Proposition 4.5, a good quality of a Θ -sketch for the sketched sparse minres approximation associated with dictionary \mathcal{D}_K on $\mathcal{P}_{\text{train}}$ can be guaranteed if Θ is an ε -embedding for every subspace $R_r(W_r; \mu)$, defined by (18), with $W_r \in \mathcal{L}_r(\mathcal{D}_K)$ and $\mu \in \mathcal{P}_{\text{train}}$. This condition can be enforced

Algorithm 3 Greedy algorithm for dictionary generation

Given: $\mathcal{P}_{\text{train}}$, $\mathbf{A}(\mu)$, $\mathbf{b}(\mu)$, $\mathbf{l}(\mu)$, Θ , τ , r .

Output: Θ -sketch of \mathbf{U}_K .

1. Set $i = 0$, $\mathcal{D}_0 = \emptyset$, obtain $\eta(\mu)$ in (45), set $\Delta^\Phi(\mu) = \|\mathbf{b}(\mu)\|_{U'}^\Phi / \eta(\mu)$ and pick $\mu^1 \in \mathcal{P}_{\text{train}}$.

while $\max_{\mu \in \mathcal{P}_{\text{train}}} \Delta^\Phi(\mathbf{u}_r(\mu); \mu) > \tau$ **do**

2. Set $i = i + 1$.

3. Evaluate $\mathbf{u}(\mu^i)$ and set $\mathcal{D}_i := \mathcal{D}_{i-1} \cup \{\mathbf{u}(\mu^i)\}$

4. Update the Θ -sketch of \mathbf{U}_i (matrix composed from the vectors in \mathcal{D}_i).

5. Use Algorithm 2 (if $i < r$, choosing $r := i$) with $\mathcal{P}_{\text{test}}$ replaced by $\mathcal{P}_{\text{train}}$ to solve (43) for all $\mu \in \mathcal{P}_{\text{train}}$.

6. Find $\mu^{i+1} := \arg \max_{\mu \in \mathcal{P}_{\text{train}}} \Delta^\Phi(\mathbf{u}_r(\mu); \mu)$.

end while

a priori for all possible outcomes of Algorithm 3 by choosing Θ such that it is an ε -embedding for every subspace $R_r(W_r; \mu)$ with $W_r \in \mathcal{L}_r(\{\mathbf{u}(\mu) : \mu \in \mathcal{P}_{\text{train}}\})$ and $\mu \in \mathcal{P}_{\text{train}}$. An embedding Θ satisfying this property with probability at least $1 - \delta$ can be obtained as a realization of a $(\varepsilon, (\#\mathcal{P}_{\text{train}})^{-1} (\#\mathcal{P}_{\text{train}})^{-1} \delta, r + 1)$ oblivious $U \rightarrow \ell_2$ subspace embedding. The computational cost of Algorithm 3 is dominated by the calculation of the snapshots $\mathbf{u}(\mu^i)$ and their Θ -sketches. As was argued in [3], the computation of the snapshots can have only a minor impact on the overall cost of an algorithm. For the classical sequential or limited-memory computational architectures, each snapshot should require a log-linear complexity and memory consumption, while for parallel and distributed computing the routines for computing the snapshots should be well-parallelizable and require low communication between cores. Moreover, for the computation of the snapshots one may use a highly-optimized commercial solver or a powerful server. The Θ -sketch of the snapshots may be computed extremely efficiently in basically any computational architecture [3, Section 4.4]. With SRHT, sketching of K snapshots requires only $\mathcal{O}(n(Km_A + m_b) \log k)$ flops, and the maintenance of the sketch requires $\mathcal{O}((Km_A + m_b)k)$ bytes of memory. By using similar arguments as in Section 4.3 it can be shown that $k = \mathcal{O}(r \log K)$ (or $k = \mathcal{O}(r^* \log K)$) is enough to yield with high probability an accurate approximation of the dictionary-based reduced model. With this value of k , the required number of flops for the computation and the amount of memory for the storage of a Θ -sketch becomes $\mathcal{O}(n(Km_A + m_b)(\log r + \log \log K))$, and $\mathcal{O}((Km_A + m_b)r \log K)$, respectively.

5 Post-processing the reduced model's solution

So far we presented a methodology for efficient computation of an approximate solution $\mathbf{u}_r(\mu)$, or to be more precise, its coordinates in a certain basis, which can be the classical reduced basis for a fixed approximation space, or the dictionary vectors for dictionary-based approximation presented in Section 4. The approximate solution $\mathbf{u}_r(\mu)$, however, is usually not what one should consider as the output. In fact, the amount of allowed online computations is highly limited and should be independent of the dimension of the full order model. Therefore outputting $\mathcal{O}(n)$ bytes of data as

$\mathbf{u}_r(\mu)$ should be avoided when $\mathbf{u}(\mu)$ is not the quantity of interest.

Further, we shall consider an approximation with a single subspace U_r noting that the presented approach can also be used for post-processing the dictionary-based approximation from Section 4 (by taking U_r as the subspace spanned by the dictionary vectors). Let \mathbf{U}_r be a matrix whose column vectors form a basis for U_r and let $\mathbf{a}_r(\mu)$ be the coordinates of $\mathbf{u}_r(\mu)$ in this basis. A general quantity of interest $s(\mu) := l(\mathbf{u}(\mu); \mu)$ can be approximated by $s_r(\mu) := l(\mathbf{u}_r(\mu); \mu)$. Further, let us assume a linear case where $l(\mathbf{u}(\mu); \mu) := \langle \mathbf{l}(\mu), \mathbf{u}(\mu) \rangle$ with $\mathbf{l}(\mu) \in U'$ being the extractor of the quantity of interest. Then

$$s_r(\mu) = \langle \mathbf{l}(\mu), \mathbf{u}_r(\mu) \rangle = \mathbf{l}_r(\mu)^H \mathbf{a}_r(\mu), \quad (47)$$

where $\mathbf{l}_r(\mu) := \mathbf{U}_r^H \mathbf{l}(\mu)$.

Remark 5.1. *In general, our approach can be used for estimating an inner product between arbitrary parameter-dependent vectors. The possible applications include efficient estimation of the primal-dual correction and an extension to quadratic quantities of interest. In particular, the estimation of the primal-dual correction can be obtained by replacing $\mathbf{l}(\mu)$ by $\mathbf{r}(\mathbf{u}_r(\mu); \mu)$ and $\mathbf{u}_r(\mu)$ by $\mathbf{v}_r(\mu)$ in (47), where $\mathbf{v}_r(\mu) \in U$ is a reduced basis (or dictionary-based) approximate solution to the adjoint problem. A quadratic output quantity of interest has the form $l(\mathbf{u}_r(\mu); \mu) := \langle \mathbf{L}(\mu) \mathbf{u}_r(\mu) + \mathbf{l}(\mu), \mathbf{u}_r(\mu) \rangle$, where $\mathbf{L}(\mu) : U \rightarrow U'$ and $\mathbf{l}(\mu) \in U'$. Such $l(\mathbf{u}_r(\mu); \mu)$ can be readily derived from (47) by replacing $\mathbf{l}(\mu)$ with $\mathbf{L}(\mu) \mathbf{u}_r(\mu) + \mathbf{l}(\mu)$.*

The affine factors of $\mathbf{l}_r(\mu)$ should be first precomputed in the offline stage and then used for online evaluation of $\mathbf{l}_r(\mu)$ for each parameter value with a computational cost independent of the dimension of the original problem. The offline computations required for evaluating the affine factors of $\mathbf{l}_r(\mu)$, however, can still be too expensive or even unfeasible to perform. Such a scenario may arise when using a high-dimensional approximation space (or a dictionary), when the extractor $\mathbf{l}(\mu)$ has many (possibly expensive to maintain) affine terms, or when working in an extreme computational environment, e.g., with data streamed or distributed among multiple workstations. In addition, evaluating $\mathbf{l}_r(\mu)$ from the affine expansion as well as evaluating $\mathbf{l}_r(\mu)^H \mathbf{a}_r(\mu)$ itself can be subject to round-off errors (especially when \mathbf{U}_r is ill-conditioned and may not be orthogonalized). Further, we shall provide a (probabilistic) way for estimating $s_r(\mu)$ with a reduced computational cost and better numerical stability. As the core we take the idea from [3, Section 4.3] proposed as a workaround to expensive offline computations for the evaluation of the primal-dual correction.

Remark 5.2. *The spaces U and U' are equipped with inner products $\langle \cdot, \cdot \rangle_U$ and $\langle \cdot, \cdot \rangle_{U'}$ (defined by matrix \mathbf{R}_U), which are used for controlling the accuracy of the approximate solution $\mathbf{u}_r(\mu)$. In general, \mathbf{R}_U is chosen according to both the operator $\mathbf{A}(\mu)$ and the extractor $\mathbf{l}(\mu)$ of the quantity of interest. The goal of this section, however, is only the estimation of the quantity of interest from the given $\mathbf{u}_r(\mu)$. Consequently, for many problems it can be more pertinent to use here a different \mathbf{R}_U than the one employed for obtaining and characterizing $\mathbf{u}_r(\mu)$. The choice for \mathbf{R}_U should be done according to $\mathbf{l}(\mu)$ (independently of $\mathbf{A}(\mu)$). For instance, for discretized parametric PDEs, if $\mathbf{l}(\mu)$ represents an integral of the solution field over the spatial domain then it is natural to choose $\langle \cdot, \cdot \rangle_U$ corresponding to the L^2 inner product. Moreover, $\langle \cdot, \cdot \rangle_U$ is required to be an inner*

product only on a certain subspace of interest, which means that \mathbf{R}_U may be a positive semi-definite matrix. This consideration can be particularly helpful when the quantity of interest depends only on the restriction of the solution field to a certain subdomain. In such a case, $\langle \cdot, \cdot \rangle_U$ can be chosen to correspond with an inner product between restrictions of functions to this subdomain. The extension of random sketching for estimation of semi-inner products is straightforward (see Remark 2.1).

5.1 Approximation of the quantity of interest

An efficiently computable and accurate estimation of $s_r(\mu)$ can be obtained in two phases. In the first phase, the manifold $\mathcal{M}_r := \{\mathbf{u}_r(\mu) : \mu \in \mathcal{P}\}$ is (accurately enough) approximated with a subspace $W_p := \text{span}(\mathbf{W}_p) \subset U$, which is spanned by an efficient to multiply (i.e., sparse or low-dimensional) matrix \mathbf{W}_p . This matrix can be selected a priori or obtained depending on \mathcal{M}_r . In Section 5.2 we shall provide some strategies for choosing or computing the columns for \mathbf{W}_p . The appropriate strategy should be selected depending on the particular problem and computational architecture. Further, the solution vector $\mathbf{u}_r(\mu)$ is approximated by its orthogonal projection $\mathbf{w}_p(\mu) := \mathbf{W}_p \mathbf{b}_p(\mu)$ on W_p . The coordinates $\mathbf{b}_p(\mu)$ can be obtained from $\mathbf{a}_r(\mu)$ by

$$\mathbf{b}_p(\mu) = \mathbf{H}_p \mathbf{a}_r(\mu), \quad (48)$$

where $\mathbf{H}_p := [\mathbf{W}_p^H \mathbf{R}_U \mathbf{W}_p]^{-1} \mathbf{W}_p^H \mathbf{R}_U \mathbf{U}_r$. Note that since \mathbf{W}_p is efficient to multiply by, the matrix \mathbf{H}_p can be efficiently precomputed in the offline stage. We arrive at the following estimation of $s_r(\mu)$:

$$s_r(\mu) \approx \langle \mathbf{l}(\mu), \mathbf{w}_p(\mu) \rangle = \mathbf{l}_r^*(\mu)^H \mathbf{a}_r(\mu), \quad (49)$$

where $\mathbf{l}_r^*(\mu)^H := \mathbf{l}(\mu)^H \mathbf{W}_p \mathbf{H}_p$. Unlike $\mathbf{l}_r(\mu)$, the affine factors of $\mathbf{l}_r^*(\mu)$ can now be efficiently pre-computed thanks to the structure of \mathbf{W}_p .

In the second phase of the algorithm, the precision of (49) is improved with a sketched (random) correction associated with an $U \rightarrow \ell_2$ subspace embedding Θ :

$$\begin{aligned} s_r(\mu) &= \langle \mathbf{l}(\mu), \mathbf{w}_p(\mu) \rangle + \langle \mathbf{l}(\mu), \mathbf{u}_r(\mu) - \mathbf{w}_p(\mu) \rangle \\ &\approx \langle \mathbf{l}(\mu), \mathbf{w}_p(\mu) \rangle + \langle \mathbf{R}_U^{-1} \mathbf{l}(\mu), \mathbf{u}_r(\mu) - \mathbf{w}_p(\mu) \rangle_U^\Theta =: s_r^*(\mu). \end{aligned} \quad (50)$$

In practice, $s_r^*(\mu)$ can be efficiently evaluated using the following relation:

$$s_r^*(\mu) = [\mathbf{l}_r^*(\mu)^H + \Delta \mathbf{l}_r^*(\mu)^H] \mathbf{a}_r(\mu), \quad (51)$$

where the affine terms of $\Delta \mathbf{l}_r^*(\mu)^H := \mathbf{l}^\Theta(\mu)^H (\mathbf{U}_r^\Theta - \mathbf{W}_p^\Theta \mathbf{H}_p)$ can be precomputed from the Θ -sketch of \mathbf{U}_r , a sketched matrix $\mathbf{W}_p^\Theta := \Theta \mathbf{W}_p$ and the matrix \mathbf{H}_p with a negligible computational cost.

Proposition 5.3. *If Θ is an $(\varepsilon, \delta, 1)$ oblivious $U \rightarrow \ell_2$ subspace embedding,*

$$|s_r(\mu) - s_r^*(\mu)| \leq \varepsilon \|\mathbf{l}(\mu)\|_{U'} \|\mathbf{u}_r(\mu) - \mathbf{w}_p(\mu)\|_U \quad (52)$$

holds for a fixed parameter $\mu \in \mathcal{P}$ with probability at least $1 - 2\delta$.

Proof. See appendix. □

Proposition 5.4. *Let $L \subset U$ denote a subspace containing $\{\mathbf{R}_U^{-1}\mathbf{l}(\mu) : \mu \in \mathcal{P}\}$. If Θ is an ε -embedding for $L + U_r + W_p$, then (52) holds for all $\mu \in \mathcal{P}$.*

Proof. See appendix. □

It follows that the accuracy of $s_r^*(\mu)$ can be controlled through the quality of W_p for approximating \mathcal{M}_r , the quality of Θ as an $U \rightarrow \ell_2$ ε -embedding, or both. Note that choosing Θ as a null matrix (i.e., an ε -embedding for U with $\varepsilon = 1$) leads to a single first-phase approximation (49), while letting $W_p := \{\mathbf{0}\}$ corresponds to a single sketched (second-phase) approximation. Such particular choices for Θ or W_p can be pertinent when the subspace W_p is highly accurate so that there is practically no benefit to use a sketched correction or, the other way around, when the computational environment or the problem does not permit a sufficiently accurate approximation of \mathcal{M}_r with W_p , therefore making the use of a non-zero $\mathbf{w}_p(\mu)$ unjustified.

Remark 5.5. *When interpreting random sketching as a Monte Carlo method for the estimation of the inner product $\langle \mathbf{l}(\mu), \mathbf{u}_r(\mu) \rangle$, the proposed approach can be interpreted as a control variate method where $\mathbf{w}_p(\mu)$ plays the role of the control variate. A multileveled Monte Carlo method with different control variates should further improve the efficiency of post-processing.*

5.2 Construction of reduced subspaces

Further we address the problem of computing the basis vectors for W_p . In general, the strategy for obtaining W_p has to be chosen according to the problem's structure and the constraints due to the computational environment.

A simple way, used in [3], is to choose W_p as the span of samples of $\mathbf{u}(\mu)$ either chosen randomly or during the first few iterations of the reduced basis (or dictionary) generation with a greedy algorithm. Such W_p , however, may be too costly to operate with. Then we propose more sophisticated constructions of W_p .

Approximate Proper Orthogonal Decomposition

A subspace W_p can be obtained by an (approximate) POD of the reduced vectors evaluated on a training set $\mathcal{P}_{\text{train}} \subseteq \mathcal{P}$. Here, randomized linear algebra can be again employed for improving efficiency. The computational cost of the proposed POD procedure shall mainly consist of the solution of $m = \#\mathcal{P}_{\text{train}}$ reduced problems and the multiplication of \mathbf{U}_r by $p = \dim(W_p) \ll r$ small vectors. Unlike the classical POD, our methodology does not require computation or maintenance of the full solution's samples and therefore allows large training sets.

Let $L_m = \{\mathbf{a}_r(\mu^i)\}_{i=1}^m$ be a training sample of the coordinates of $\mathbf{u}_r(\mu)$ in a basis \mathbf{U}_r . We look for a POD subspace W_r associated with the snapshot matrix

$$\mathbf{W}_m := [\mathbf{u}_r(\mu^1), \mathbf{u}_r(\mu^2), \dots, \mathbf{u}_r(\mu^m)] = \mathbf{U}_r \mathbf{L}_m,$$

where \mathbf{L}_m is a matrix whose columns are the elements from L_m .

An accurate estimation of POD can be efficiently computed via the sketched method of snapshots introduced in [3, Section 5.2]. More specifically, a quasi-optimal (with high probability) POD basis can be calculated as

$$\mathbf{W}_p := \mathbf{U}_r \mathbf{T}_p^*, \quad (53)$$

where

$$\mathbf{T}_p^* := \mathbf{L}_m[\mathbf{t}_1, \dots, \mathbf{t}_p],$$

with $\mathbf{t}_1, \dots, \mathbf{t}_p$ being the p dominant singular vectors of $\mathbf{U}_r^\Theta \mathbf{L}_m$. Note that the matrix \mathbf{T}_p^* can be efficiently obtained with a computational cost independent of the dimension of the full order model. The dominant cost is the multiplication of \mathbf{U}_r by \mathbf{T}_p^* , which is also expected to be inexpensive since \mathbf{T}_p^* has a small number of columns. Guarantees for the quasi-optimality of W_p can be readily derived from [3, Theorem 5.5].

Sketched greedy algorithm

A greedy search over the training set $\{\mathbf{u}_r(\mu) : \mu \in \mathcal{P}_{\text{train}}\}$ of approximate solutions is another way to construct W_p . At the i -th iteration, \mathbf{W}_i is enriched with a vector $\mathbf{u}_r(\mu^{i+1})$ with the largest distance to W_i over the training set. Note that in this case the resulting matrix \mathbf{W}_p has the form (53), where $\mathbf{T}_p^* = [\mathbf{a}_r(\mu^1), \dots, \mathbf{a}_r(\mu^p)]$. The efficiency of the greedy selection can be improved by employing random sketching technique. At each iteration, the distance to W_i can be measured with the sketched norm $\|\cdot\|_U^\Theta$, which can be computed from sketches $\Theta \mathbf{u}_r(\mu) = \mathbf{U}_r^\Theta \mathbf{a}_r(\mu)$ of the approximate solutions with no need to operate with large matrix \mathbf{U}_r but only its sketch. This allows efficient computation of the quasi-optimal interpolation points μ^1, \dots, μ^p and the associated matrix \mathbf{T}_p^* . Note that for numerical stability an orthogonalization of \mathbf{W}_i with respect to $\langle \cdot, \cdot \rangle_U^\Theta$ can be performed, that can be done by modifying \mathbf{T}_i^* so that $\mathbf{U}_r^\Theta \mathbf{T}_i^*$ is an orthogonal matrix. Such \mathbf{T}_i^* can be obtained with standard QR factorization. When \mathbf{T}_p^* has been computed, the matrix \mathbf{W}_p can be calculated by multiplying \mathbf{U}_r with \mathbf{T}_p^* . The quasi-optimality of μ^1, \dots, μ^p and approximate orthogonality of \mathbf{W}_p is guaranteed if Θ is an ε -embedding for all subspaces from the set $\{W_p + \text{span}(\mathbf{u}_r(\mu^i))\}_{i=1}^m$. This property of Θ can be guaranteed a priori by considering $(\varepsilon, \binom{m}{p}^{-1} \delta, p+1)$ oblivious $U \rightarrow \ell_2$ subspace embeddings, or certified a posteriori with the procedure explained in Section 6.

Coarse approximation

Let us notice that the online cost of evaluating $s_r^*(\mu)$ does not depend on the dimension p of W_p . Consequently, if W_p is spanned by structured (e.g., sparse) basis vectors then a rather high dimension is allowed (possibly larger than r).

For classical numerical methods for PDEs, the resolution of the mesh (or grid) is usually chosen to guarantee both an approximability of the solution manifold by the approximation space and the stability. For many problems the latter factor is dominant and one chooses the mesh primary to it. This is a typical situation for wave problems, advection-diffusion-reaction problems and many others. For these problems, the resolution of the mesh can be much higher than needed for the

estimation of the quantity of interest from the given solution field. In these cases, the quantity of interest can be efficiently yet accurately approximated using a coarse-grid interpolation of the solution.

Suppose that each vector $\mathbf{u} \in U$ represents a function $u : \Omega \rightarrow \mathbb{K}$ in a finite-dimensional approximation space spanned by basis functions $\{\psi_i(x)\}_{i=1}^n$ associated with a fine mesh of Ω . The function $u(x)$ can be approximated by a projection on a coarse-grid approximation space spanned by basis functions $\{\phi_i(x)\}_{i=1}^p$. For simplicity assume that each $\phi_i(x) \in \text{span}\{\psi_j(x)\}_{j=1}^n$. Then the i -th basis vector for W_p can be obtained simply by evaluating the coordinates of $\phi_i(x)$ in the basis $\{\psi_j(x)\}_{j=1}^n$. Note that for the classical finite element approximation, each basis function has a local support and the resulting matrix \mathbf{W}_p is sparse.

6 A posteriori certification of the sketch and solution

Here we provide a simple, yet efficient procedure for a posteriori verification of the quality of a sketching matrix and describe a few scenarios where such a procedure can be employed. The proposed a posteriori certification of the sketched reduced model and its solution is probabilistic. It does not require operating with high-dimensional vectors but only with their small sketches. The quality of a certificate shall be characterized by two user specified parameters: $0 < \delta^* < 1$ for the probability of success and $0 < \varepsilon^* < 1$ for the tightness of the computed error bounds.

6.1 Verification of an ε -embedding for a given subspace

Let Θ be a $U \rightarrow \ell_2$ subspace embedding and $V \subset U$ be a subspace of U (chosen depending on the reduced model, e.g., $V := R_r(U_r; \mu)$ in (18)). Recall that the quality of Θ can be measured by the accuracy of $\langle \cdot, \cdot \rangle_U^\Theta$ as an approximation of $\langle \cdot, \cdot \rangle_U$ for vectors in V .

We propose to verify the accuracy of $\langle \cdot, \cdot \rangle_U^\Theta$ simply by comparing it to an inner product $\langle \cdot, \cdot \rangle_U^{\Theta^*}$ associated with a new random embedding $\Theta^* \in \mathbb{K}^{k^* \times n}$, where Θ^* is chosen such that for any given vectors $\mathbf{x}, \mathbf{y} \in V$ the concentration inequality

$$|\langle \mathbf{x}, \mathbf{y} \rangle_U - \langle \mathbf{x}, \mathbf{y} \rangle_U^{\Theta^*}| \leq \varepsilon^* \|\mathbf{x}\|_U \|\mathbf{y}\|_U \quad (54)$$

holds with probability at least $1 - \delta^*$. One way to ensure (54) is to choose Θ^* as an $(\varepsilon^*, \delta^*, 1)$ oblivious $U \rightarrow \ell_2$ subspace embedding. A condition on the number of rows for the oblivious embedding can be either obtained theoretically (see [3]) or chosen from the practical experience, which should be the case for embeddings constructed with SRHT matrices (recall that they have worse theoretical guarantees than the Gaussian matrices but perform equally well in practice). Alternatively, Θ^* can be built by random sampling of the rows of a larger ε -embedding for V . This approach can be far more efficient than generating Θ^* as an oblivious embedding (see Remark 6.1) or even essential for some computational architectures. Another requirement for Θ^* is that it is generated independently from Θ . Therefore, in the algorithms we suggest to consider Θ^* only for the certification of the solution and nothing else.

Remark 6.1. *In some scenarios it can be beneficial to construct Θ and Θ^* by sampling their rows from a fixed realization of a larger oblivious embedding $\hat{\Theta}$, which is guaranteed a priori to be an ε -embedding for V with high probability. More precisely, Θ and Θ^* can be defined as*

$$\Theta := \Gamma \hat{\Theta}, \quad \Theta^* := \Gamma^* \hat{\Theta}, \quad (55)$$

where Γ and Γ^* are random independent sampling (or Gaussian, or SRHT) matrices. In this way, a $\hat{\Theta}$ -sketch of a reduced order model can be first precomputed and then used for efficient evaluation/update of the sketches associated with Θ and Θ^* . This approach can be essential for the adaptive selection of the optimal size for Θ in a limited-memory environment where only one pass (or a few passes) over the reduced basis vectors is allowed and therefore there is no chance to recompute a sketch associated with an oblivious embedding at each iteration. It may also reduce the complexity (number of flops) of an algorithm (especially when Θ is constructed with SRHT matrices) by not requiring to recompute high-dimensional matrix-vector products multiple times.

Let \mathbf{V} denote a matrix whose columns form a basis of V . Define the sketches $\mathbf{V}^\Theta := \Theta \mathbf{V}$ and $\mathbf{V}^{\Theta^*} := \Theta^* \mathbf{V}$. Note that \mathbf{V}^Θ and \mathbf{V}^{Θ^*} contain as columns low-dimensional vectors and therefore are cheap to maintain and to operate with (unlike the matrix \mathbf{V}).

We start with the certification of the inner product between two fixed vectors from V (see Proposition 6.2).

Proposition 6.2. *For any two vectors $\mathbf{x}, \mathbf{y} \in V$ possibly depending on Θ but independent of Θ^* , we have that*

$$\begin{aligned} |\langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{U}}^{\Theta^*} - \langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{U}}^{\Theta}| - \frac{\varepsilon^*}{1 - \varepsilon^*} \|\mathbf{x}\|_{\mathcal{U}}^{\Theta^*} \|\mathbf{y}\|_{\mathcal{U}}^{\Theta^*} &\leq |\langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{U}} - \langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{U}}^{\Theta}| \\ &\leq |\langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{U}}^{\Theta^*} - \langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{U}}^{\Theta}| + \frac{\varepsilon^*}{1 - \varepsilon^*} \|\mathbf{x}\|_{\mathcal{U}}^{\Theta^*} \|\mathbf{y}\|_{\mathcal{U}}^{\Theta^*} \end{aligned} \quad (56)$$

holds with probability at least $1 - 4\delta^*$.

Proof. See appendix. □

The error bounds in Proposition 6.2 can be computed from the sketches of \mathbf{x} and \mathbf{y} , which may be efficiently evaluated from \mathbf{V}^Θ and \mathbf{V}^{Θ^*} and the coordinates of \mathbf{x} and \mathbf{y} associated with \mathbf{V} , with no operations on high-dimensional vectors. A certification for several pairs of vectors should be obtained using a union bound for the probability of success. By replacing \mathbf{x} by $\mathbf{R}_{\mathcal{U}}^{-1} \mathbf{x}'$ and \mathbf{y} by $\mathbf{R}_{\mathcal{U}}^{-1} \mathbf{y}'$ in Proposition 6.2 and using definition (2) one can derive a certification of the dual inner product $\langle \cdot, \cdot \rangle_{\mathcal{U}}^{\Theta}$ for vectors \mathbf{x}' and \mathbf{y}' in $V' := \{\mathbf{R}_{\mathcal{U}} \mathbf{x} : \mathbf{x} \in V\}$.

In general, the quality of an approximation with a Θ -sketch of a reduced model should be characterized by the accuracy of $\langle \cdot, \cdot \rangle_{\mathcal{U}}^{\Theta}$ for the whole subspace V . Let ω be the minimal value for ε such that Θ satisfies an ε -embedding property for V . Now, we address the problem of computing an a posteriori upper bound $\bar{\omega}$ for ω from the sketches \mathbf{V}^Θ and \mathbf{V}^{Θ^*} and we provide conditions to ensure quasi-optimality of $\bar{\omega}$.

Proposition 6.3. *For a fixed realization of Θ^* , let us define*

$$\bar{\omega} := \max \left\{ 1 - (1 - \varepsilon^*) \min_{\mathbf{x} \in V/\{\mathbf{0}\}} \left(\frac{\|\mathbf{x}\|_{\Theta}}{\|\mathbf{x}\|_{\Theta^*}} \right)^2, (1 + \varepsilon^*) \max_{\mathbf{x} \in V/\{\mathbf{0}\}} \left(\frac{\|\mathbf{x}\|_{\Theta}}{\|\mathbf{x}\|_{\Theta^*}} \right)^2 - 1 \right\}. \quad (57)$$

If $\bar{\omega} < 1$, then Θ is guaranteed to be a $U \rightarrow \ell_2$ $\bar{\omega}$ -subspace embedding for V with probability at least $1 - \delta^$.*

Proof. See appendix. □

It follows that if $\bar{\omega} < 1$ then it is an upper bound for ω with high probability. Assume that \mathbf{V}^{Θ} and \mathbf{V}^{Θ^*} have full ranks. Let \mathbf{T}^* be the matrix such that $\mathbf{V}^{\Theta^*} \mathbf{T}^*$ is orthogonal (with respect to ℓ_2 -inner product). Such a matrix can be computed with a QR factorization. Then $\bar{\omega}$ defined in (57) can be obtained from the following relation

$$\bar{\omega} = \max \left\{ 1 - (1 - \varepsilon^*) \sigma_{\min}^2, (1 + \varepsilon^*) \sigma_{\max}^2 - 1 \right\}, \quad (58)$$

where σ_{\min} and σ_{\max} are the minimal and the maximal singular values of the small matrix $\mathbf{V}^{\Theta} \mathbf{T}^*$.

We have that $\bar{\omega} \geq \varepsilon^*$. The value for ε^* may be selected an order of magnitude less than ω with no considerable impact on the computational cost, therefore in practice the effect of ε^* on $\bar{\omega}$ can be considered to be negligible. Proposition 6.3 implies that $\bar{\omega}$ is an upper bound of ω with high probability. A guarantee of effectivity of $\bar{\omega}$ (i.e., its closeness to ω), however, has not been yet provided. To do so we shall need a stronger assumption on Θ^* than (54).

Proposition 6.4. *If the realization of Θ^* is a $U \rightarrow \ell_2$ ω^* -subspace embedding for V , then $\bar{\omega}$ (defined by (57)) satisfies*

$$\bar{\omega} \leq \frac{1 + \varepsilon^*}{1 - \omega^*} (1 + \omega) - 1. \quad (59)$$

Proof. See appendix. □

If Θ^* is a $(\omega^*, \gamma^*, \dim(V))$ oblivious $U \rightarrow \ell_2$ subspace embedding, then the condition on Θ^* in Proposition 6.4 is satisfied with probability at least $1 - \gamma^*$ (for some user-specified value γ^*). Therefore, a matrix Θ^* of moderate size should yield a very good upper bound $\bar{\omega}$ of ω . Moreover, if Θ and Θ^* are drawn from the same distribution, then Θ^* can be expected to be an ω^* -embedding for V with $\omega^* = \mathcal{O}(\omega)$ with high probability. Combining this consideration with Proposition 6.4 we deduce that a sharp upper bound should be obtained for some $k^* \leq k$. Therefore, in the algorithms one may readily consider $k^* := k$. If pertinent, a better value for k^* can be selected adaptively, at each iteration increasing k^* by a constant factor until the desired tolerance or a stagnation of $\bar{\omega}$ is reached.

6.2 Certification of a sketch of a reduced model and its solution

The results of Propositions 6.2 and 6.3 can be employed for certification of a sketch of a reduced model and its solution. They can also be used for adaptive selection of the number of rows of a random sketching matrix to yield an accurate approximation of the reduced model. Thereafter we discuss several practical applications of the methodology described above.

Approximate solution

Let $\mathbf{u}_r(\mu) \in U$ be an approximation of $\mathbf{u}(\mu)$. The accuracy of $\mathbf{u}_r(\mu)$ can be measured with the residual error $\|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'}$, which can be efficiently estimated by

$$\|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'} \approx \|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'}^{\Theta}.$$

The certification of such estimation can be derived from Proposition 6.2 choosing $\mathbf{x} = \mathbf{y} := \mathbf{R}_U^{-1}\mathbf{r}(\mathbf{u}_r(\mu); \mu)$ and using definition (2) of $\|\cdot\|_{U'}^{\Theta}$.

For applications, which involve computation of snapshots over the training set (e.g., approximate POD or greedy algorithm with the exact error indicator), one should be able to efficiently precompute the sketches of $\mathbf{u}(\mu)$. Then the error $\|\mathbf{u}(\mu) - \mathbf{u}_r(\mu)\|_U$ can be efficiently estimated by

$$\|\mathbf{u}(\mu) - \mathbf{u}_r(\mu)\|_U \approx \|\mathbf{u}(\mu) - \mathbf{u}_r(\mu)\|_U^{\Theta}.$$

Such an estimation can be certified with Proposition 6.2 choosing $\mathbf{x} = \mathbf{y} := \mathbf{u}(\mu) - \mathbf{u}_r(\mu)$.

Output quantity

In Section 5 we provided a way for estimating the output quantity $s_r(\mu) = \langle \mathbf{l}(\mu), \mathbf{u}_r(\mu) \rangle$. More specifically $s_r(\mu)$ can be efficiently estimated by $s_r^*(\mu)$ defined in (51). We have

$$|s_r(\mu) - s_r^*(\mu)| = |\langle \mathbf{R}_U^{-1}\mathbf{l}(\mu), \mathbf{u}_r(\mu) - \mathbf{w}_p(\mu) \rangle_U - \langle \mathbf{R}_U^{-1}\mathbf{l}(\mu), \mathbf{u}_r(\mu) - \mathbf{w}_p(\mu) \rangle_U^{\Theta}|,$$

therefore the quality of $s_r^*(\mu)$ may be certified by Proposition 6.2 with $\mathbf{x} = \mathbf{R}_U^{-1}\mathbf{l}(\mu)$, $\mathbf{y} = \mathbf{u}_r(\mu) - \mathbf{w}_p(\mu)$.

Minimal residual projection

By Proposition 3.3, the quality of the Θ -sketch of a subspace U_r for approximating the minres projection for a given parameter value can be characterized by the lowest value ω for ε such that Θ satisfies the ε -embedding property for subspace $V := R_r(U_r; \mu)$, defined in (18). The upper bound $\bar{\omega}$ of such ω can be efficiently computed using (57). The verification of a Θ -sketch for all parameter values in \mathcal{P} , simultaneously, can be performed by considering a subspace V in (57), which contains $\bigcup_{\mu \in \mathcal{P}} R_r(U_r; \mu)$.

Dictionary-based approximation

For each parameter value, the quality of Θ for dictionary-based approximation defined in (43) can be characterized by the quality of the sketched minres projection associated with a subspace $U_r(\mu)$, which can be verified by computing $\bar{\omega}$ in (57) associated with $V := R_r(U_r(\mu); \mu)$.

Adaptive selection of the size for a random sketching matrix

When no a priori bound for the size of Θ sufficient to yield an accurate sketch of a reduced model is available, or when the bounds are pessimistic, the sketching matrix should be selected adaptively. At each iteration, if the certificate indicates a poor quality of a Θ -sketch for approximating the solution (or the error) on $\mathcal{P}_{\text{train}} \subseteq \mathcal{P}$, one can improve the accuracy of the sketch by adding extra rows to Θ . In the analysis, the embedding Θ^* used for certification was assumed to be independent of Θ , consequently a new realization of Θ^* should be sampled after each decision to improve Θ has been made. To save computational costs, the previous realizations of Θ^* and the associated Θ^* -sketches can be readily recycled as parts of the updates for Θ and the Θ -sketch.

We finish with a practical application of Propositions 6.3 and 6.4. Consider a situation where one is given a class of random embeddings (e.g., oblivious subspace embeddings mentioned in Section 2.1 or the embeddings constructed with random sampling of rows of an ε -embedding as in Remark 6.1) and one is interested in generating an ε -embedding Θ (or rather computing the associated sketch), with a user-specified accuracy $\varepsilon \leq \tau$, for V (e.g., a subspace containing $\cup_{\mu \in \mathcal{P}} R_r(U_r; \mu)$) with nearly optimal number of rows. Moreover, we consider that no bound of the size of matrices to yield an ε -embedding is available or that the bound is pessimistic. It is only known that matrices with more than k_0 rows satisfy (54). This condition could be derived theoretically (as for Gaussian matrices) or deduced from practical experience (for SRHT). The matrix Θ can be readily generated adaptively using $\bar{\omega}$ defined by (57) as an error indicator (see Algorithm 4). It directly follows by a union bound argument that Θ generated in Algorithm 4 is an ε -embedding for V , with $\varepsilon \leq \tau$, with probability at least $1 - t\delta^*$, where t is the number of iterations taken by the algorithm.

Algorithm 4 Adaptive selection of the number of rows for Θ

Given: $k_0, \mathbf{V}, \tau > \varepsilon^*$.

Output: $\mathbf{V}^\Theta, \mathbf{V}^{\Theta^*}$

1. Set $k = k_0$ and $\bar{\omega} = \infty$.

while $\bar{\omega} > \tau$ **do**

2. Generate Θ and Θ^* with k rows and evaluate $\mathbf{V}^\Theta := \Theta\mathbf{V}$ and $\mathbf{V}^{\Theta^*} := \Theta^*\mathbf{V}$.

3. Use (58) to compute $\bar{\omega}$.

4. Increase k by a constant factor.

end while

To improve the efficiency, at each iteration of Algorithm 4 we could select the number of rows for Θ^* adaptively instead of choosing it equal to k . In addition, the embeddings from previous iterations can be considered as parts of Θ at further iterations.

7 Numerical experiments

This section is devoted to experimental validation of the methodology as well as realization of its great practical potential. The numerical tests were carried out on two benchmark problems that are

difficult to tackle with the standard projection-based MOR methods due to a high computational cost and issues with numerical stability of the computation (or minimization) of the residual norm, or bad approximability of the solution manifold with a low-dimensional space.

In all the experiments we used oblivious $U \rightarrow \ell_2$ embeddings of the form

$$\Theta := \Omega \mathbf{Q},$$

where \mathbf{Q} was taken as the (sparse) transposed Cholesky factor of \mathbf{R}_U (or another norm-inducing matrix as in Remark 5.2) and Ω as a SRHT matrix. The random embedding Γ used for the online efficiency was also taken as SRHT. Moreover, for simplicity in all the experiments the coefficient $\eta(\mu)$ for the error estimation was chosen as 1.

The experiments were executed on an Intel[®] Core[™] i7-7700HQ 2.8GHz CPU, with 16.0GB RAM memory using Matlab[®] R2017b.

7.1 Acoustic invisibility cloak

The first numerical example is inspired by the development of invisibility cloaking [10, 11]. It consists in an acoustic wave scattering in 2D with a perfect scatterer covered in an invisibility cloak composed of layers of homogeneous isotropic materials. The geometry of the problem is depicted in Figure 1a. The cloak consists of 32 layers of equal thickness 1.5625 cm each constructed with 4 sublayers of equal thickness of alternating materials: mercury (a heavy liquid) followed by a light liquid. The properties (density and bulk modulus) of the light liquids are chosen to minimize the visibility of the scatterer for the frequency band [7.5, 8.5] kHz. The associated boundary value problem with first order absorbing boundary conditions is the following

$$\begin{cases} \nabla \cdot (\rho^{-1} \nabla u) + \rho^{-1} \kappa^2 u = 0, & \text{in } \Omega \\ (j\kappa - \frac{1}{2R_\Omega})u + \frac{\partial u}{\partial n} = (j\kappa - \frac{1}{2R_\Omega})u^{in} + \frac{\partial u^{in}}{\partial n}, & \text{on } \Gamma \\ \frac{\partial u}{\partial n} = 0, & \text{on } \Gamma_s, \end{cases} \quad (60)$$

where $j = \sqrt{-1}$, $u = u^{in} + u^{sc}$ is the total pressure, with $u^{in} = \exp(-j\kappa(y-4))$ Pa · m being the pressure of the incident plane wave and u^{sc} being the pressure of the scattered wave, ρ is the material's density, $\kappa = \frac{2\pi f}{c}$ is the wave number, $c = \sqrt{\frac{b}{\rho}}$ is the speed of sound and b is the bulk modulus. The background material is chosen as water having density $\rho = \rho_0 = 997$ kg/m³ and bulk modulus $b = b_0 = 2.23$ GPa. For the frequency $f = 8$ kHz the associated wave number of the background is $\kappa = \kappa_0 = 33.6$ m⁻¹. The i -th layer of the cloak (enumerated starting from the outermost layer) is composed of 4 alternating layers of mercury with density $\rho = \rho_m = 13500$ kg/m³ and bulk modulus $b = b_m = 28$ GPa and a light liquid with density $\rho = \rho_i$ and bulk modulus $b = b_i$ given in Table 1. The light liquids from Table 1 can in practice be obtained, for instance, with the pentamode mechanical metamaterials [5, 22].

The last 10 layers contain liquids with small densities that can be subject to imperfections during the manufacturing process. Moreover, the external conditions (such as temperature and pressure) may also affect the material's properties. We then consider a characterization of the impact of small perturbations of the density and the bulk modulus of the light liquids in the last

Table 1: The properties, density in kg/m³ and bulk modulus in GPa, of the light liquids in the cloak.

i	ρ_i	b_i	i	ρ_i	b_i	i	ρ_i	b_i	i	ρ_i	b_i
1	231	0.483	9	179	0.73	17	56.1	0.65	25	9	1.34
2	121	0.328	10	166	0.78	18	59.6	0.687	26	9	2.49
3	162	0.454	11	150	0.745	19	40.8	0.597	27	9	2.5
4	253	0.736	12	140	0.802	20	32.1	0.682	28	9	2.5
5	259	0.767	13	135	0.786	21	22.5	0.521	29	9	0.58
6	189	0.707	14	111	0.798	22	15.3	0.6	30	9.5	1.91
7	246	0.796	15	107	0.8	23	10	0.552	31	9.31	0.709
8	178	0.739	16	78	0.656	24	9	1.076	32	9	2.44

10 layers on the quality of the cloak in the frequency regime [7.8, 8.2] kHz. Assuming that the density and the bulk modulus may vary by 2.5%, the corresponding parameter set is

$$\mathcal{P} = \prod_{23 \leq i \leq 32} [0.975\rho_i, 1.025\rho_i] \times \prod_{23 \leq i \leq 32} [0.975b_i, 1.025b_i] \times [7.8 \text{ kHz}, 8.2 \text{ kHz}].$$

Note that in this case $\mathcal{P} \subset \mathbb{R}^{21}$.

The quantity of interest is chosen to be the following

$$s(\mu) = l(u(\mu); \mu) \|u(\mu) - u^{in}(\mu)\|_{L^2(\Omega_1)}^2 / b_0 = \|u^{sc}(\mu)\|_{L^2(\Omega_1)}^2 / b_0,$$

which represents the (rescaled, time-averaged) acoustic energy of the scattered wave concealed in the region Ω_1 (see Figure 1a). For the considered parameter set $s(\mu)$ is ranging from $0.0225A_s$ to $0.095A_s$, where $A_s = \|u^{in}\|_{L^2(\Omega_1)}^2 / b_0 = 7.2 \text{ J} \cdot \text{Pa} / b_0$ at frequency 8 kHz.

The problem is symmetric with respect to the $x = 0$ axis, therefore only half of the domain has to be considered for discretization. For the discretization, we used piecewise quadratic approximation on a mesh of triangular (finite) elements. The mesh was chosen such that there were at least 20 degrees of freedom per wavelength, which is a standard choice for Helmholtz problems with a moderate wave number. It yielded approximately 400000 complex degrees of freedom for the discretization. Figures 1b to 1d depict the solutions $u(\mu)$ for different parameter values with quantities of interest $s(\mu) = 0.033A_s, 0.044A_s$ and $0.065A_s$, respectively.

It is revealed that for this problem, considering the classical H^1 inner product for the solution space leads to dramatic instabilities of the projection-based MOR methods. To improve the stability, the inner product is chosen corresponding to the specific structure of the operator in (60). The solution space U is here equipped with the following inner product

$$\langle \mathbf{v}, \mathbf{w} \rangle_U := \langle \rho_s^{-1} \kappa_s^2 v, w \rangle_{L^2} + \langle \rho_s^{-1} \nabla v, \nabla w \rangle_{L^2}, \quad \mathbf{v}, \mathbf{w} \in U, \quad (61)$$

where v and w are the functions identified with \mathbf{v} and \mathbf{w} , respectively, and ρ_s and κ_s are the density and the wave number associated with the unperturbed cloak (i.e., with properties from Table 1) at frequency 8 kHz.

The operator for this benchmark is directly given in an affine form with $m_A = 23$ terms. Furthermore, for online efficiency we used EIM to obtain an approximate affine representation of

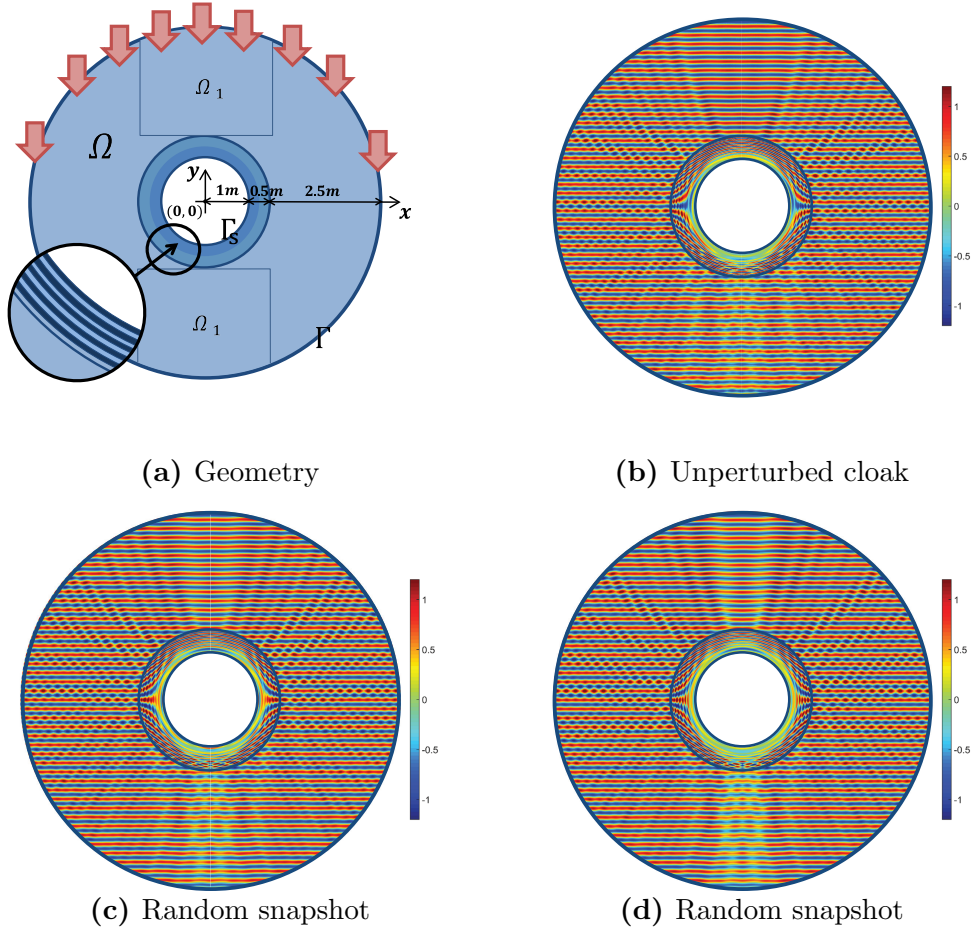


Figure 1: (a) Geometry of the invisibility cloak benchmark. (b) The real component of u in $\text{Pa} \cdot \text{m}$ for the parameter value $\mu \in \mathcal{P}$ corresponding to Table 1 and frequency $f = 8$ kHz. (c)-(d) The real component of u in $\text{Pa} \cdot \text{m}$ for two random samples from \mathcal{P} with $f = 7.8$ kHz.

$u^{in}(\mu)$ (or rather a vector $\mathbf{u}^{in}(\mu)$ from U representing an approximation of $u^{in}(\mu)$) and the right hand side vector with 50 affine terms (with error close to machine precision). The approximation space U_r of dimension $r = 150$ was constructed with a greedy algorithm (based on sketched minres projection) performed on a training set of 50000 uniform random samples in \mathcal{P} . The test set $\mathcal{P}_{test} \subset \mathcal{P}$ was taken as 1000 uniform random samples in \mathcal{P} .

Minimal residual projection. Let us first address the validation of the sketched minres projection from Section 3.2. For this we computed sketched (and standard) minres projections $\mathbf{u}_r(\mu)$ of $\mathbf{u}(\mu)$ onto U_r for each $\mu \in \mathcal{P}_{test}$ with sketching matrix Θ of varying sizes. The error of approximation is here characterized by $\Delta_{\mathcal{P}} := \max_{\mu \in \mathcal{P}_{test}} \|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'} / \|\mathbf{b}(\mu_s)\|_{U'}$ and $e_{\mathcal{P}} := \max_{\mu \in \mathcal{P}_{test}} \|\mathbf{u}(\mu) - \mathbf{u}_r(\mu)\|_U / \|\mathbf{u}^{in}(\mu_s)\|_U$, where $\mathbf{u}^{in}(\mu_s)$ is the vector representing the incident wave and $\mathbf{b}(\mu_s)$ is the right hand side vector associated with the unperturbed cloak and the

frequency $f = 8$ kHz (see Figures 2a and 2c). Furthermore, in Figure 2e we provide the characterization of the maximal error in the quantity of interest $e_{\mathcal{P}}^s := \max_{\mu \in \mathcal{P}_{\text{test}}} |s(\mu) - s_r(\mu)|/A_s$. For each size of Θ , 20 realizations of the sketching matrix were considered to analyze the statistical properties of $e_{\mathcal{P}}$, $\Delta_{\mathcal{P}}$ and $e_{\mathcal{P}}^s$.

For comparison, along with the minimal residual projections we also computed the sketched (and classical) Galerkin projection introduced in [3]. Figures 2b, 2d and 2f depict the errors $\Delta_{\mathcal{P}}$, $e_{\mathcal{P}}$, $e_{\mathcal{P}}^s$ of a sketched (and classical) Galerkin projection using Θ of different sizes. Again, for each k we used 20 realizations of Θ to characterize the statistical properties of the error. We see that the classical Galerkin projection is more accurate in the exact norm $\|\cdot\|_U$ and the quantity of interest than the standard minres projection. On the other hand, it is revealed that the minres projection is far better suited to random sketching.

From Figure 2 one can clearly report the (essential) preservation of the quality of the classical minres projection for $k \geq 500$. Note that for the minres projection a small deviation of $e_{\mathcal{P}}$ and $e_{\mathcal{P}}^s$ is observed. These errors are higher or lower than the standard values with (almost) equal probability (for $k \geq 500$). In contrast to the minres projection, the quality of the Galerkin projection is not preserved even for large k up to 10000. This can be explained by the fact that the approximation of the Galerkin projection with random sketching is highly sensitive to the properties of the operator, which here is non-coercive and has a high condition number (for some parameter values), while the (essential) preservation of the accuracy of the standard minres projection by its sketched version is guaranteed regardless of the operator's properties. One can clearly see that the sketched minres projection using Θ with just $k = 500$ rows yields better approximation (in terms of the maximal observed error) of the solution than the sketched Galerkin projection with $k = 5000$, even though the standard minres projection is less accurate than the Galerkin one.

As already discussed, random sketching improves not only efficiency but also has an advantage of making the reduced model less sensitive to round-off errors thanks to possibility of direct online solution of the (sketched) least-squares problem without appealing to the normal equation. Figure 3 depicts the maximal condition number $\kappa_{\mathcal{P}}$ over $\mathcal{P}_{\text{test}}$ of the reduced matrix $\mathbf{V}_r^{\Theta}(\mu) := \Theta \mathbf{R}_U^{-1} \mathbf{A}(\mu) \mathbf{U}_r$ associated with the sketched minres projection using reduced basis matrix \mathbf{U}_r with (approximately) unit-orthogonal columns with respect to $\langle \cdot, \cdot \rangle_U$, for varying sizes of Θ . We also provide the maximal condition number of the reduced (normal) system of equations associated with the classical minres projection. It is observed that indeed random sketching yields an improvement of numerical stability by a square root.

Approximation of the output quantity. The next experiment was performed for a fixed approximation $\mathbf{u}_r(\mu)$ obtained with sketched minres projection using Θ with 1000 rows. For such $\mathbf{u}_r(\mu)$, the approximate extraction of the quantity $s_r(\mu) = l(\mathbf{u}_r(\mu); \mu)$ from $\mathbf{u}_r(\mu)$ (represented by coordinates in reduced basis) was considered with the efficient procedure from Section 5.

The post-processing procedure was performed by choosing $\mathbf{l}(\mu)$ and $\mathbf{u}_r(\mu)$ in (47) as $\mathbf{u}_r(\mu) - \mathbf{u}^{in}(\mu)$. Furthermore, for better accuracy the solution space was here equipped with (semi-)inner product $\langle \cdot, \cdot \rangle_U := \langle \cdot, \cdot \rangle_{L^2(\Omega_1)}$ that is different from the inner product (61) considered for ensuring quasi-optimality of the minres projection and error estimation (see Remark 5.2). For such choices of $\mathbf{l}(\mu)$, $\mathbf{u}_r(\mu)$ and $\langle \cdot, \cdot \rangle_U$, we employed a greedy search with error indicator $\min_{\mathbf{w} \in W_i} \|\mathbf{u}_r(\mu) - \mathbf{w}\|_U^{\Theta}$ over training set of 50000 uniform samples in \mathcal{P} to find W_p . Then $s_r(\mu)$ was efficiently approximated

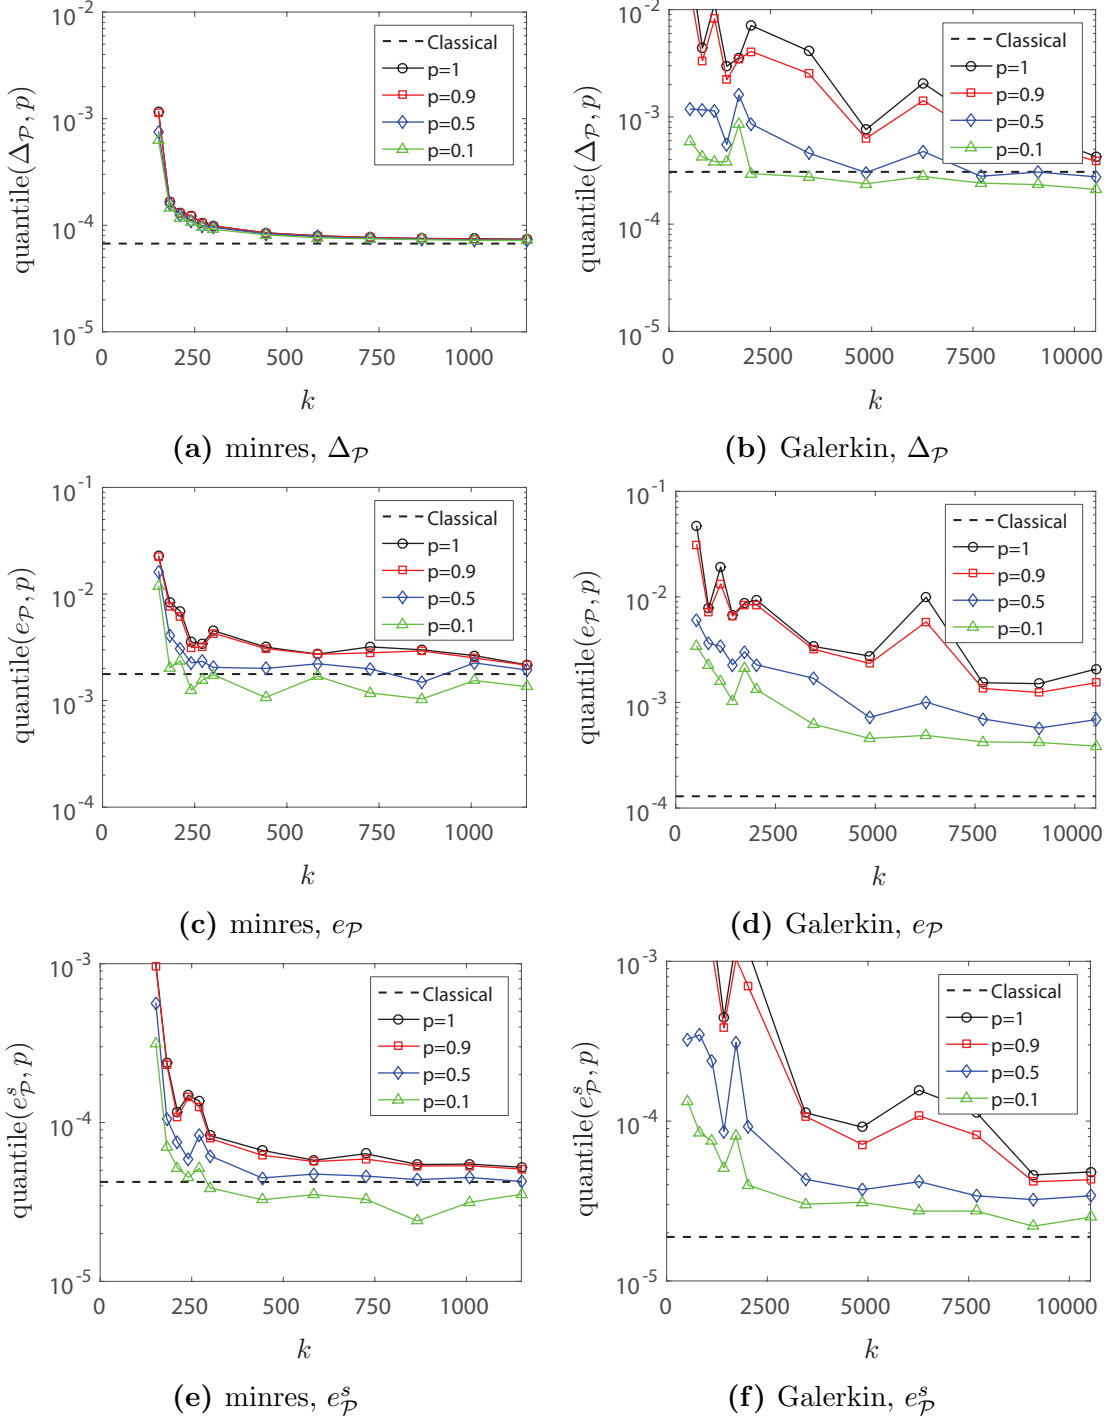


Figure 2: The errors of the classical minres and Galerkin projections and quantiles of probabilities $p = 1, 0.9, 0.5$ and 0.1 over 20 realizations of the errors of the sketched minres and Galerkin projections, versus the number of rows of Θ . (a) Residual error $\Delta_{\mathcal{P}}$ of standard and sketched minres projection. (b) Residual error $\Delta_{\mathcal{P}}$ of standard and sketched Galerkin projection. (c) Exact error $e_{\mathcal{P}}$ (in $\|\cdot\|_U$) of standard and sketched minres projection. (d) Exact error $e_{\mathcal{P}}$ (in $\|\cdot\|_U$) of standard and sketched Galerkin projection. (e) Quantity of interest error $e_{\mathcal{P}}^s$ of standard and sketched minres projection. (f) Quantity of interest error $e_{\mathcal{P}}^s$ of standard and sketched Galerkin projection.




Figure 3: The maximal condition number over $\mathcal{P}_{\text{test}}$ of the reduced (normal) system associated with the classical minres projection and quantiles of probabilities $p = 1, 0.9, 0.5$ and 0.1 over 20 realizations of the maximal condition number of the sketched reduced matrix $\mathbf{V}_r^\Theta(\mu)$, versus the number of rows k of Θ .

by $s_r^*(\mu)$ given in (51). In this experiment, the error is characterized by $e_{\mathcal{P}}^s = \max_{\mu \in \mathcal{P}_{\text{test}}} |s(\mu) - \tilde{s}_r(\mu)|/A_s$, where $\tilde{s}_r(\mu) = s_r(\mu)$ or $s_r^*(\mu)$. The statistical properties of $e_{\mathcal{P}}^s$ for each value of k and $\dim(W_p)$ were obtained with 20 realizations of Θ . Figure 4a exhibits the dependence of $e_{\mathcal{P}}^s$ on the size of Θ with W_p of dimension $\dim(W_p) = 15$. Furthermore, in Figure 4b we provide the maximum value $e_{\mathcal{P}}^s$ from the computed samples for different sizes of Θ and W_p . The accuracy of $s_r^*(\mu)$ can be controlled by the quality of W_p for the approximation of $\mathbf{u}_r(\mu)$ and the quality of $\langle \cdot, \cdot \rangle_U^\Theta$ for the approximation of $\langle \cdot, \cdot \rangle_U$. When W_p approximates well $\mathbf{u}_r(\mu)$, one can use a random correction with Θ of rather small size, while in the alternative scenario the usage of a large random sketch is required. In this experiment we see that the quality of the output is nearly preserved with high probability when using W_p of dimension $\dim(W_p) = 20$ and a sketch of size $k = 1000$, or W_p of dimension $\dim(W_p) = 15$ and a sketch of size $k = 10000$. For less accurate W_p , with $\dim(W_p) \leq 10$, the preservation of the quality of the output requires larger sketches of sizes $k \geq 30000$. For optimizing efficiency the dimension for W_p and the size for Θ should be picked depending on the dimensions r and n of U_r and U , respectively, and the particular computational architecture. The increase of the considered dimension of W_p entails storage and operation with more high-dimensional vectors, while the increase of the sketch entails higher computational cost associated with storage and operation with the sketched matrix $\mathbf{U}_r^\Theta = \Theta \mathbf{U}_r$. Let us finally note that for this benchmark the approximate extraction of the quantity of interest with the procedure from Section 5 using $\dim(W_p) = 15$ and a sketch of size $k = 10000$, required in about 10 times less amount of storage and complexity than the classical exact extraction.

Certification of the sketch. Next the experimental validation of the procedure for a posteriori certification of the Θ -sketch or the sketched solution (see Section 6) is addressed. For this, we




Figure 4: The error $e_{\mathcal{P}}^s$ of $s_r(\mu)$ or its efficient approximation $s_r^*(\mu)$ using W_p and Θ of varying sizes. (a) The error of $s_r(\mu)$ and quantiles of probabilities $p = 1, 0.9, 0.5$ and 0.1 over 20 realizations of $e_{\mathcal{P}}^s$ associated with $s_r^*(\mu)$ using W_p with $\text{dim}(W_p) = 10$, versus sketch's size k . (b) The error of $s_r(\mu)$ and maximum over 20 realizations of $e_{\mathcal{P}}^s$ associated with $s_r^*(\mu)$, versus sketch's size k for W_p of varying dimension.

generated several Θ of different sizes k and for each of them computed the sketched minres projections $\mathbf{u}_r(\mu) \in U_r$ for all $\mu \in \mathcal{P}_{\text{test}}$. Thereafter Propositions 6.2 and 6.3, with $V(\mu) := R_r(U_r; \mu)$ defined by (18), were considered for certification of the residual error estimates $\|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'}^{\Theta}$ or the quasi-optimality of $\mathbf{u}_r(\mu)$ in the residual error. Oblivious embeddings of varying sizes were tested for Θ^* . For simplicity it was assumed that all considered Θ^* satisfy (54) with $\varepsilon^* = 0.05$ and (small) probability of failure δ^* .

By Proposition 6.2 the certification of the sketched residual error estimator $\|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'}^{\Theta}$ can be performed by comparing it to $\|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'}^{\Theta^*}$. More specifically, by (56) we have that with probability at least $1 - 4\delta^*$,

$$\begin{aligned} & \left| \|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'}^2 - (\|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'}^{\Theta})^2 \right|^{1/2} \\ & \leq \left(\left| (\|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'}^{\Theta^*})^2 - (\|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'}^{\Theta})^2 \right| + \frac{\varepsilon^*}{1 - \varepsilon^*} (\|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'}^{\Theta^*})^2 \right)^{1/2}. \end{aligned} \quad (62)$$

Figure 5 depicts $d_{\mathcal{P}} := \max_{\mu \in \mathcal{P}_{\text{test}}} d(\mathbf{u}_r(\mu); \mu) / \|\mathbf{b}(\mu_s)\|_{U'}$, where $d(\mathbf{u}_r(\mu); \mu)$ is the exact discrepancy $\|\|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'}^2 - (\|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'}^{\Theta})^2\|^{1/2}$ or its (probabilistic) upper bound in (62). For each Θ and k^* , 20 realizations of $d_{\mathcal{P}}$ were computed for statistical analysis. We see that (sufficiently) tight upper bounds for $\|\|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'}^2 - (\|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'}^{\Theta})^2\|^{1/2}$ were obtained already when $k^* \geq 100$, which is in particular several times smaller than the size of Θ required for quasi-optimality of $\mathbf{u}_r(\mu)$. This implies that the certification of the effectivity of the error estimator $\|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'}^{\Theta}$ by $\|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'}^{\Theta^*}$ should require negligible computational costs compared to the cost of obtaining the solution (or estimating the error in adaptive algorithms such as greedy algorithms).

By Proposition 3.3, the quasi-optimality of $\mathbf{u}_r(\mu)$ can be guaranteed if Θ is an ε -embedding for $V(\mu)$. The ε -embedding property of each Θ was verified with Proposition 6.3. In Figure 6 we

provide $\omega_{\mathcal{P}} := \max_{\mu \in \mathcal{P}_{\text{test}}} \tilde{\omega}(\mu)$ where $\tilde{\omega}(\mu) = \omega(\mu)$, which is the minimal value for ε such that Θ is an ε -embedding for $V(\mu)$, or $\tilde{\omega}(\mu) = \bar{\omega}(\mu)$, which is the upper bound of $\omega(\mu)$ computed with (58) using Θ^* of varying sizes. For illustration purposes we here allow the value ε in Definition 2.2 to be larger than 1. The statistical properties of $\omega_{\mathcal{P}}$ were obtained with 20 realizations for each Θ and value of k^* . Figure 6a depicts the statistical characterization of $\omega_{\mathcal{P}}$ for Θ of size $k = 5000$. The maximal value of $\omega_{\mathcal{P}}$ observed for each k^* and Θ is presented in Figure 6b. It is observed that with a posteriori estimates from Proposition 6.3 using Θ^* of size $k^* = 6000$, we here can guarantee with high probability that Θ with $k = 5000$ satisfies an ε -embedding property for $\varepsilon \approx 0.6$. The theoretical bounds from [3] for Θ to be an ε -embedding for $V(\mu)$ with $\varepsilon = 0.6$ yield much larger sizes, namely, for the probability of failure $\delta \leq 10^{-6}$, they require more than $k = 45700$ rows for Gaussian matrices and $k = 102900$ rows for SRHT. This proves Proposition 6.3 to be very useful for the adaptive selection of sizes of random matrices or for the certification of the sketched inner product $\langle \cdot, \cdot \rangle_U^{\Theta}$ for all vectors in V . Note that the adaptive selection of the size of Θ can also be performed without requiring Θ to be an ε -embedding for $V(\mu)$ with $\varepsilon < 1$, based on the observation that oblivious embeddings yield preservation of the quality of the minres projection when they are ε -embeddings for $V(\mu)$ with small ε , which is possibly larger than 1 (see Remark 7.1).

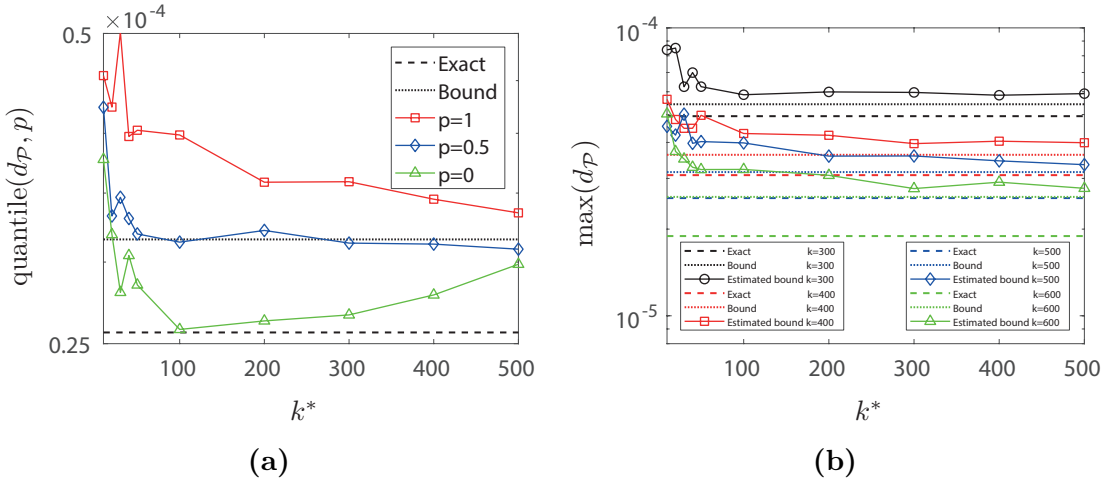


Figure 5: The discrepancy $|\|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'}^2 - (\|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'}^{\Theta})^2|^{1/2}$ between the residual error and the sketched error estimator with Θ , and the upper bound of this value computed with (62). (a) The exact discrepancy for Θ with $k = 500$ rows, the upper bound (62) of this discrepancy taking $\|\cdot\|_{U'}^{\Theta^*} = \|\cdot\|_{U'}$, and quantiles of probabilities $p = 1, 0.5$ and 0 (i.e., the observed maximum, median and minimum) over 20 realizations of the (probabilistic) upper bound (62) versus the size of Θ^* . (b) The exact discrepancy, the upper bound (62) taking $\|\cdot\|_{U'}^{\Theta^*} = \|\cdot\|_{U'}$, and the maximum of 20 realizations of the (probabilistic) upper bound (62) versus the number of rows k^* of Θ^* for varying sizes k of Θ .

Remark 7.1. Throughout the paper the quality of Θ (e.g., for approximation of minres projection in Section 3.2) was characterized by ε -embedding property. However, for this numerical benchmark the sufficient size for Θ to be an ε -embedding for $V(\mu)$ is in several times larger than the one




Figure 6: The minimal value for ε such that Θ is an ε -embedding for $V_r(\mu)$ for all $\mu \in \mathcal{P}_{\text{test}}$, and a posteriori random estimator of this value obtained with the procedure from Section 6 using Θ^* with k^* rows. (a) The minimal value for ε for Θ with $k = 5000$ rows and quantiles of probabilities $p = 1, 0.9, 0.5$ and 0.1 over 20 samples of the estimator, versus the size of Θ^* . (b) The minimal value for ε and the maximum of 20 samples of the estimator, versus the number of rows of Θ^* for varying sizes of Θ .

yielding an accurate approximation of the minres projection. In particular, Θ with $k = 500$ rows with high probability provides an approximation with residual error very close to the minimal one, but it does not satisfy an ε -embedding property (with $\varepsilon < 1$), which is required for guaranteeing the quasi-optimality of $\mathbf{u}_r(\mu)$ with Proposition 3.3. A more reliable way for certification of the quality of Θ for approximation of the minres projection onto U_r can be derived by taking into account that Θ was generated from a distribution of oblivious embeddings. In such a case it is enough to only certify that $\|\cdot\|_U^\Theta$ provides an approximate upper bound of $\|\cdot\|_U$ for all vectors in $V(\mu)$ without the need to guarantee that $\|\cdot\|_U^\Theta$ is an approximate lower bound (that in practice is the main bottleneck). This approach is outlined below.

We first observe that Θ was generated from a distribution of random matrices such that for all $\mathbf{x} \in V(\mu)$, we have

$$\mathbb{P} \left(\left| \|\mathbf{x}\|_U^2 - (\|\mathbf{x}\|_U^\Theta)^2 \right| \leq \varepsilon_0 \|\mathbf{x}\|_U^2 \right) \geq 1 - \delta_0.$$

The values ε_0 and δ_0 can be obtained from the theoretical bounds from [3] or practical experience. Then one can show that for the sketched minres projection $\mathbf{u}_r(\mu)$ associated with Θ , the relation

$$\|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'} \leq \sqrt{\frac{1 + \varepsilon_0}{1 - \omega(\mu)}} \min_{\mathbf{x} \in U_r} \|\mathbf{r}(\mathbf{x}; \mu)\|_{U'}, \quad (63)$$

holds with probability at least $1 - \delta_0$, where $\omega(\mu) < 1$ is the minimal value for ε such that for all $\mathbf{x} \in V(\mu)$

$$(1 - \varepsilon) \|\mathbf{x}\|_U^2 \leq (\|\mathbf{x}\|_U^\Theta)^2.$$

The quasi-optimality of $\mathbf{u}_r(\mu)$ in the norm $\|\cdot\|_U$ rather than the residual norm can be readily derived from relation (63) by using the equivalence between the residual norm and the error in $\|\cdot\|_U$.

In this way a characterization of the quasi-optimality of the sketched minres projection with Θ can be obtained from the a posteriori upper bound of $\omega(\mu)$ in (63). Note that since Θ is an oblivious subspace embedding, the parameters ε_0 and δ_0 do not depend on the dimension of $V(\mu)$, which implies that the considered value for ε_0 should be an order of magnitude less than $\omega(\mu)$. Therefore, it can be a good way to choose ε_0 as $\omega(\mu)$ (or rather its upper bound) multiplied by a small factor, say 0.1.

The (probabilistic) upper bound $\bar{\omega}(\mu)$ for $\omega(\mu)$ can be obtained a posteriori by following a similar procedure as the one from Proposition 6.3 described for verification of the ε -embedding property. More precisely, we can use similar arguments as in Proposition 6.3 to show that

$$\bar{\omega}(\mu) := 1 - (1 - \varepsilon^*) \min_{\mathbf{x} \in V/\{\mathbf{0}\}} \left(\frac{\|\mathbf{x}\|_U^\Theta}{\|\mathbf{x}\|_U^{\Theta^*}} \right)^2$$

is an upper bound for $\omega(\mu)$ with probability at least $1 - \delta^*$.

Let us now provide experimental validation of the proposed approach. For this we considered same sketching matrices Θ as in the previous experiment for validation of the ε -embedding property. For each Θ we computed $\omega_{\mathcal{P}} := \max_{\mu \in \mathcal{P}_{\text{test}}} \tilde{\omega}(\mu)$, where $\tilde{\omega}(\mu) = \omega(\mu)$ or its upper bound $\bar{\omega}(\mu)$ using Θ^* of different sizes (see Figure 7). Again 20 realizations of $\omega_{\mathcal{P}}$ were considered for the statistical characterization of $\omega_{\mathcal{P}}$ for each Θ and size of Θ^* . One can clearly see that the present approach provides better estimation of the quasi-optimality constants than the one with the ε -embedding property. In particular, the quasi-optimality guarantee for Θ with $k = 500$ rows is experimentally verified. Furthermore, we see that in all the experiments the a posteriori estimates are lower than 1 even for Θ^* of small sizes, yet they are larger than the exact values, which implies efficiency and robustness of the method. From Figure 7, a good accuracy of a posteriori estimates is with high probability attained for $k^* \geq k/2$.

Computational costs. For this benchmark, random sketching yielded drastic computational savings in the offline stage and considerably improved online efficiency. To verify the gains for the offline stage, we executed two greedy algorithms for the generation of the reduced approximation space of dimension $r = 150$ based on the minres projection and the sketched minres projection, respectively. The standard algorithm resulted in a computational burden after reaching 96-th iteration due to exceeding the limit of RAM (16GB). Note that performing $r = 150$ iterations in this case would require around 25GB of RAM (mainly utilized for storage of the affine factors of $\mathbf{R}_U^{-1} \mathbf{A}(\mu) \mathbf{U}_r$). In contrast to the standard method, conducting $r = 150$ iterations of a greedy algorithm with random sketching using Θ of size $k = 2500$ (and $\mathbf{\Gamma}$ of size $k' = 500$) for the sketched minres projection and Θ^* of size $k^* = 250$ for the error certification, took only 0.65GB of RAM. Moreover, the sketch required only a minor part (0.2GB) of the aforementioned amount of memory, while the major part was consumed by the initialization and maintenance of the full order model. The sketched greedy algorithm had a total runtime of 1.9 hours, from which 0.8 hours was spent on the computation of 150 snapshots, 0.2 hours on the provisional online solutions and 0.9 hours on random projections. Note that a drastic reduction of the offline runtime would as well be observed

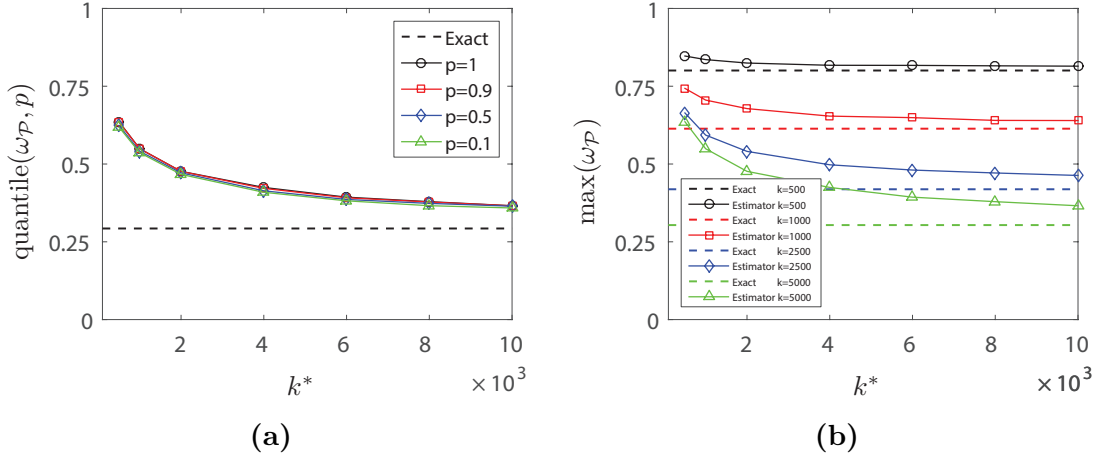


Figure 7: The minimal value for ε such that $(1 - \varepsilon)\|\mathbf{x}\|_U^2 \leq (\|\mathbf{x}\|_U^\Theta)^2$ holds for all $\mathbf{x} \in V_r(\mu)$ and $\mu \in \mathcal{P}_{\text{test}}$, and a posteriori random estimator of this value using Θ^* with k^* rows. (a) The minimal value for ε for Θ with $k = 5000$ rows and quantiles of probabilities $p = 1, 0.9, 0.5$ and 0.1 over 20 realizations of the estimator, versus the size of Θ^* . (b) The minimal value for ε and the maximum of 20 realizations of the estimator versus the number of rows of Θ^* , for varying sizes of Θ .

even in a computational environment with higher RAM (than 25GB), since random sketching in addition to reducing the memory consumption, also greatly improves the efficiency in terms of complexity and other metrics.

Next the improvement of online computational cost of minres projection is addressed. For this, we computed the reduced solutions on the test set with a standard method, which consists in assembling the reduced system of equations (representing the normal equation) from its affine decomposition (precomputed in the offline stage) and its subsequent solution with built in Matlab[®] R2017b linear solver. The online solutions on the test set were additionally computed with the sketched method for comparison of runtimes and storage requirements. For this, for each parameter value, the reduced least-squares problem was assembled from the precomputed affine decompositions of $\mathbf{V}_r^\Phi(\mu)$ and $\mathbf{b}^\Phi(\mu)$ and solved with the normal equation using the built in Matlab[®] R2017b linear solver. Note that both methods proceeded with the normal equation. The difference was in the way how this equation was obtained. For the standard method it was directly assembled from the affine representation, while for the sketched method it was computed from the sketched matrices $\mathbf{V}_r^\Phi(\mu)$ and $\mathbf{b}^\Phi(\mu)$.

Table 2 depicts the runtimes and memory consumption taken by the standard and sketched online stages for varying sizes of the reduced space and Φ (for the sketched method). The sketch's sizes were picked such that the associated reduced solutions with high probability had almost (higher by at most a factor of 1.2) optimal residual error. Our approach nearly halved the online runtime for all values of r from Table 2. Furthermore, the improvement of memory requirements was even greater. For instance, for $r = 150$ the online memory consumption was divided by 6.8.

Table 2: CPU times in seconds and amount of memory in MB taken by the standard and the efficient sketched online solvers for the solutions on the test set.

	Standard			Sketched		
	$r = 50$	$r = 100$	$r = 150$	$r = 50$ $k' = 300$	$r = 100$ $k' = 400$	$r = 150$ $k' = 500$
CPU	1.6	5.5	12	0.9	2.8	5.3
Storage	22	87	193	5.8	15	28

7.2 Advection-diffusion problem

The dictionary-based approximation method proposed in Section 4 is validated on a 2D advection dominated advection-diffusion problem defined on a complex flow. This problem is governed by the following equations

$$\begin{cases} -\epsilon\Delta u + \boldsymbol{\beta} \cdot \nabla u = f, & \text{in } \Omega \\ u = 0, & \text{on } \Gamma_{out} \\ \frac{\partial u}{\partial \mathbf{n}} = 0, & \text{on } \Gamma_n, \end{cases} \quad (64)$$

where u is the unknown (temperature) field, $\epsilon = 0.0001$ is the diffusion (heat conduction) coefficient and $\boldsymbol{\beta}$ is the advection field. The geometry of the problem is as follows. First we have 5 circular pores of radius 0.01 located at points $\mathbf{x}_j = 0.5(\cos(2\pi j/5), \sin(2\pi j/5))$, $1 \leq j \leq 5$. The domain of interest is then defined as the square $[-10, 10]^2$ without the pores, i.e, $\Omega := [-10, 10]^2/\Omega_n$, with $\Omega_n := \cup_{1 \leq j \leq 5} \{\mathbf{x} : \|\mathbf{x} - \mathbf{x}_j\| \leq 0.01\}$. The boundaries Γ_n and Γ_{out} are taken as $\partial\Omega_n$ and $\partial\Omega/\partial\Omega_n$, respectively. Furthermore, Ω is (notationally) divided into the main region inside $[-1, 1]^2$, and the outer domain playing a role of a boundary layer. Finally, the force term f is nonzero in the disc $\Omega_s := \{\mathbf{x} \in \Omega : \|\mathbf{x}\| \leq 0.025\}$. The geometric setup of the problem is presented in Figure 8a.

The advection field is taken as a potential (divergence-free and curl-free) field consisting of a linear combination of 12 components,


$$\boldsymbol{\beta}(\mathbf{x}) = \mu_2 \cos(\mu_1)\hat{\mathbf{e}}_x + \mu_2 \sin(\mu_1)\hat{\mathbf{e}}_y + \sum_{i=1}^{10} \mu_i \boldsymbol{\beta}_i(\mathbf{x}), \quad \mathbf{x} \in \Omega,$$

where


$$\boldsymbol{\beta}_i(\mathbf{x}) = \begin{cases} \frac{-\hat{\mathbf{e}}_r(\mathbf{x}_i)}{\|\mathbf{x} - \mathbf{x}_i\|} & \text{for } 1 \leq i \leq 5 \\ \frac{-\hat{\mathbf{e}}_\theta(\mathbf{x}_{i-5})}{\|\mathbf{x} - \mathbf{x}_{i-5}\|} & \text{for } 6 \leq i \leq 10. \end{cases} \quad (65)$$

The vectors $\hat{\mathbf{e}}_x$ and $\hat{\mathbf{e}}_y$ are the basis vectors of the Cartesian system of coordinates. The vectors $\hat{\mathbf{e}}_r(\mathbf{x}_j)$ and $\hat{\mathbf{e}}_\theta(\mathbf{x}_j)$ are the basis vectors of the polar coordinate system with the origin at point \mathbf{x}_j , $1 \leq j \leq 5$. From the physics perspective, we have here a superposition of two uniform flows and five hurricane flows (each consisting of a sink and a rotational flow) centered at different locations. The source term is


$$f(\mathbf{x}) = \begin{cases} \frac{1}{\pi 0.025^2} & \text{for } \mathbf{x} \in \Omega_s, \\ 0 & \text{for } \mathbf{x} \in \Omega/\Omega_s. \end{cases}$$




(a) Geometry



(b) Snapshot at μ_s



(c) Random snapshot



(d) Random snapshot

Figure 8: (a) Geometry of the advection-diffusion problem. (b) The solution field u for parameter value $\mu_s := (0, 0, 0.308, 0.308, 0.308, 0.308, 0.308, 0.616, 0.616, 0.616, 0.616, 0.616)$. (c)-(d) The solution field u for two random samples from \mathcal{P} .

We consider a multi-objective scenario, where one aims to approximate the average solution field $s^j(u)$, $1 \leq j \leq 15$, inside sensor Ω_j having a form of a disc of radius 0.025 located as in Figure 8a. The objective is to obtain sensor outputs for the parameter values $\mu := (\mu_1, \dots, \mu_{12}) \in \mathcal{P} := \{[0, 2\pi] \times [0, 0.028] \times [0.308, 0.37]^5 \times [0.616, 0.678]^5\}$. Figures 8a to 8c present solutions $u(\mu)$ for few samples from \mathcal{P} .

The discretization of the problem was performed with the classical finite element method. A nonuniform mesh was considered with finer elements near the pores of the hurricanes, and larger ones far from the pores such that each element's Peclet number inside $[-1, 1]^2$ was larger than 1 for any parameter value in \mathcal{P} . Moreover, it was revealed that for this benchmark the solution field outside the region $[-1, 1]^2$ was practically equal to zero for all $\mu \in \mathcal{P}$. Therefore the outer region was discretized with coarse elements. For the discretization we used about 380000 and 20000 degrees of freedom in the main region and the outside boundary layer, respectively, which yielded approximately 400000 degrees of freedom in total.

The solution space is equipped with the inner product

$$\|\mathbf{w}\|_U^2 := \|\nabla w\|_{L^2}^2, \quad \mathbf{w} \in U,$$

which is compatible with the H_0^1 inner product.

For this problem, approximation of the solution with a fixed low-dimensional space is ineffective. The problem has to be approached with non-linear approximation methods with parameter-dependent approximation spaces. For this, the classical hp -refinement method is computationally intractable due to high dimensionality of the parameter domain, which makes the dictionary-based approximation to be the most pertinent choice.

The training and test sets $\mathcal{P}_{\text{train}}$ and $\mathcal{P}_{\text{test}}$ were respectively chosen as 20000 and 1000 uniform random samples from \mathcal{P} . Then, Algorithm 3 was employed to generate dictionaries of sizes $K = 1500$, $K = 2000$ and $K = 2500$ for the dictionary-based approximation with $r = 100$, $r = 75$ and $r = 50$ vectors, respectively. For comparison, we also performed a greedy reduced basis algorithm (based on sketched minres projection) to generate a fixed reduced approximation space, which in particular coincides with Algorithm 3 with large enough r (here $r = 750$). Moreover, for more efficiency (to reduce the number of online solutions) at i -th iteration of Algorithm 3 and reduced basis algorithm instead of taking μ^{i+1} as a maximizer of $\Delta^\Phi(\mathbf{u}_r(\mu); \mu)$ over $\mathcal{P}_{\text{train}}$, we relaxed the problem to finding any parameter-value such that

$$\Delta^\Phi(\mathbf{u}_r(\mu^{i+1}); \mu^{i+1}) \geq \max_{\mu \in \mathcal{P}_{\text{train}}} \min_{1 \leq j \leq i} \Delta^\Phi(\mathbf{u}_r^j(\mu); \mu), \quad (66)$$

where $\mathbf{u}_r^j(\mu)$ denotes the solution obtained at the j -th iteration. Note that (66) improved the efficiency, yet yielding at least as accurate maximizer of the dictionary-based width (defined in (25)) as considering $\mu^{i+1} := \arg \max_{\mu \in \mathcal{P}_{\text{train}}} \Delta^\Phi(\mathbf{u}_r(\mu); \mu)$. For the error certification purposes, each 250 iterations the solution was computed on the whole training set and μ^{i+1} was taken as $\arg \max_{\mu \in \mathcal{P}_{\text{train}}} \Delta^\Phi(\mathbf{u}_r(\mu); \mu)$. Figure 9 depicts the observed evolutions of the errors in the greedy algorithms.

We see that at the first r iterations, the error decay for the dictionary generation practically coincides with the error decay of the reduced basis algorithm, which can be explained by the fact




Figure 9: Evolutions of the errors in Algorithm 3 for the dictionary generation for varying values of r , and the reduced basis greedy algorithm based on (sketched) minres projection. (a) The residual error $\tilde{\Delta}(\mu^{i+1}) := \|\mathbf{r}(\mathbf{u}_r(\mu^{i+1}); \mu^{i+1})\|_{U'} / \|\mathbf{b}\|_{U'}$. (b) The minimal value of the error at parameter value μ^{i+1} at the first i iterations.

that the first r iterations of the two algorithms coincide. The convergence rates remain the same for the reduced basis algorithm (even at high iterations), while they slowly subsequently degrade for dictionary-based approximation. The later method still highly outperforms the former one, since its online computational cost scales only linearly with the number of iterations. Furthermore, for the dictionary-based approximation the convergence of the error is moderately noisy. The noise is primarily due to approximating the solutions of online sparse least-squares problems with the orthogonal greedy algorithm, for which the accuracy can be sensitive to the enrichment of the dictionary with new vectors. The quality of online solutions could be improved by the usage of more sophisticated methods for sparse least-squares problems.

As it is clear from Figure 9, the obtained dictionaries and the bases taken from them provide approximations at least as accurate (on the training set) as the minres approximation with a fixed reduced space of dimension $r = 750$. Yet, the dictionary-based approximations are much more online-efficient. Table 3 provides the online complexity and storage requirements for obtaining the dictionary-based solutions for all $\mu \in \mathcal{P}_{\text{test}}$ (recall, $\#\mathcal{P}_{\text{test}} = 1000$) with the orthogonal greedy algorithm (Algorithm 2) from a sketch of size $k = 8r$, and the sketched minres solutions with QR factorization with Householder transformations (as in Matlab[®] R2017b least-squares solver) of the sketched reduced matrix in (21) from a sketch of size $k = 4r$. In particular, we see that the dictionary-based approximation with $r = 75$ and $K = 2000$ yields a gain in complexity by a factor of 15 and memory consumption by a factor of 1.9. In Table 3 we also provide the associated runtimes and required RAM. It is revealed that the dictionary-based approximation with $K = 2000$ and $r = 75$ had an about 4 times speedup. The difference between the gains in terms of complexity and runtime can be explained by high efficiency of the Matlab[®] R2017b least-squares solver. It is important to note that even more considerable enhance of efficiency could be obtained by better exploitation of the structure of the dictionary-based reduced model, in particular, by representing

the sketched matrix $\mathbf{V}_K^\Theta(\mu)$ in a format well suited for the orthogonal greedy algorithm (e.g., a product of a dense matrix by several sparse matrices similarly as in [25, 29]).

Table 3: Computational cost of obtaining online solutions for all parameter values from the test set with the reduced basis method (based on sketched minres projection) and the dictionary-based approximations.

	RB, $r = 750$	$K = 1500, r = 100$	$K = 2000, r = 75$	$K = 2500, r = 50$
Complexity in flops	3.1×10^9	0.27×10^9	0.2×10^9	0.12×10^9
Storage in flns	2.9×10^8	1.6×10^8	1.6×10^8	1.3×10^8
CPU in s	400	124	113	100
Storage in MB	234	124	124	104

Further we provide statistical analysis of the dictionary-based approximation with $K = 2000$ and $r = 75$. For this we computed the associated dictionary-based solutions $\mathbf{u}_r(\mu)$ for all parameter values in the test set, considering Θ of varying sizes. The accuracy of an approximation is characterized by the quantities $\Delta_{\mathcal{P}} := \max_{\mu \in \mathcal{P}_{\text{test}}} \|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'} / \|\mathbf{b}\|_{U'}$, $e_{\mathcal{P}} := \max_{\mu \in \mathcal{P}_{\text{test}}} \|\mathbf{u}(\mu) - \mathbf{u}_r(\mu)\|_U / \max_{\mu \in \mathcal{P}_{\text{test}}} \|\mathbf{u}(\mu)\|_U$ and $e_{\mathcal{P}}^i = \max_{\mu \in \mathcal{P}_{\text{test}}} |s^i(\mathbf{u}(\mu)) - s^i(\mathbf{u}_r(\mu))|$, $1 \leq i \leq 15$. Figure 10 depicts the dependence of $\Delta_{\mathcal{P}}$, $e_{\mathcal{P}}$ and $e_{\mathcal{P}}^i$ (for few selected values of i) on the number of rows k of Θ . For each value of k , the statistical properties of $\Delta_{\mathcal{P}}$, $e_{\mathcal{P}}$ and $e_{\mathcal{P}}^i$ were characterized with 20 realizations of $\Delta_{\mathcal{P}}$, $e_{\mathcal{P}}$ and $e_{\mathcal{P}}^i$. It is observed that for $k = 600$, the errors $\Delta_{\mathcal{P}}$ and $e_{\mathcal{P}}$ are concentrated around 0.03 and 0.06, respectively. Moreover, for all tested $k \geq 600$ we obtained nearly the same errors, which suggests preservation of the quality of the dictionary-based approximation by its sketched version at $k = 600$. A (moderate) deviation of the errors in the quantities of interest (even for very large k) can be explained by (moderately) low effectivity of representation of these errors with the error in $\|\cdot\|_U$, which we considered to control.

8 Conclusion

In this article we have extended the methodology from [3] to minres methods and proposed a novel nonlinear approximation method to tackle problems with a slow decay of Kolmogorov r -width. Furthermore, we additionally proposed efficient randomized methods for extraction of the quantity of interest and a posteriori certification of the reduced model’s sketch. The results from this paper can be used as a remedy of the drawbacks revealed in [3].

First, a method to approximate minres projection via random sketching was introduced. For each parameter value, the approximation is obtained by minimization of the ℓ_2 -norm of a random projection of the residual. The associated minimization problem can be assembled from a sketch of a low-dimensional space containing the residual vectors and then solved with a standard routine such as QR factorization or (less stable) direct solution of the normal equation. This procedure enables drastic reduction of the computational cost in any modern computational architecture and improvement of numerical stability of the standard method. Precise conditions on the sketch to yield the (approximate) preservation of the quasi-optimality constants of the standard minres

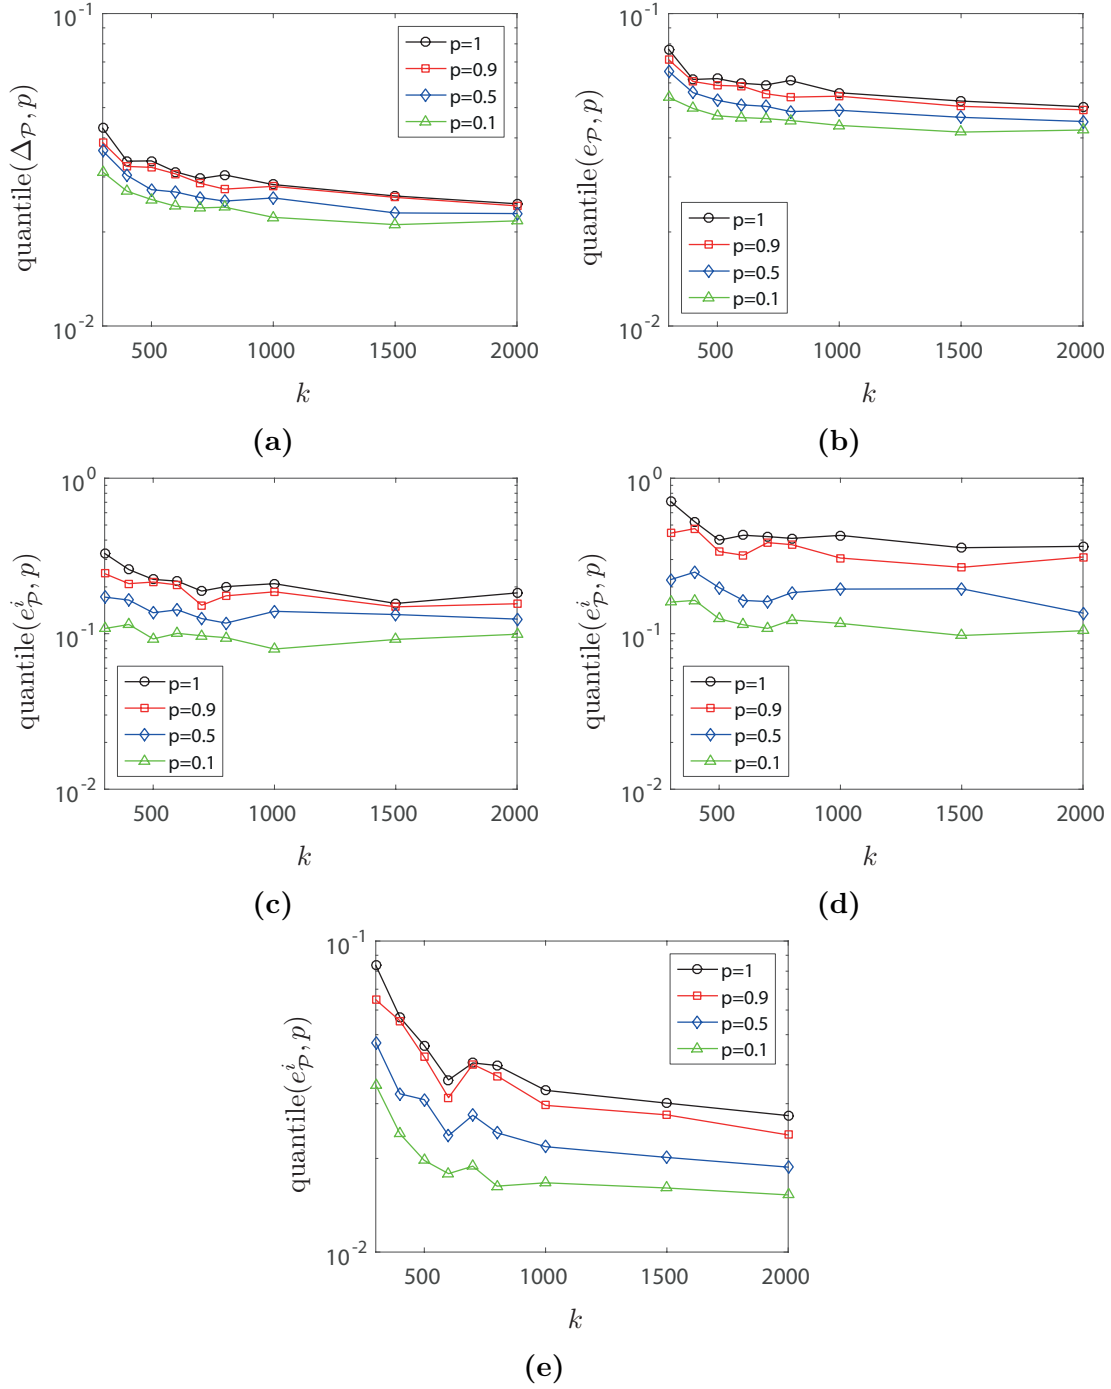


Figure 10: Quantiles of probabilities $p = 1, 0.9, 0.5$ and 0.1 over 20 samples of the errors $\Delta_{\mathcal{P}}$, $e_{\mathcal{P}}$, $e_{\mathcal{P}}^i$ of the dictionary-based approximation with $K = 2000$ and $r = 75$, versus the number of rows of Θ . (a) Residual error $\Delta_{\mathcal{P}}$. (b) Exact error $e_{\mathcal{P}}$. (c) Error $e_{\mathcal{P}}^i$ in the quantity of interest associated with sensor $i = 13$. (d) Error $e_{\mathcal{P}}^i$ in the quantity of interest associated with sensor $i = 8$. (e) Error $e_{\mathcal{P}}^i$ in the quantity of interest associated with sensor $i = 1$.

projection were provided. The conditions do not depend on the operator’s properties, which implies robustness for ill-conditioned and non-coercive problems.

Then we proposed a dictionary-based approximation, where the solution is approximated by a minres projection onto a parameter-dependent reduced space with basis vectors (adaptively) selected from a dictionary. The characterization of the quasi-optimality of the proposed dictionary-based minres projection was provided. For each parameter value the solution can be efficiently approximated by the online solution of a small sparse least-squares problem assembled from random sketches of the dictionary vectors, which entails practical feasibility of the method. It was further shown that the preservation of the quasi-optimality constants of the sparse minres projection by its sketched version can be guaranteed if the random projection satisfies an ε -embedding property for a collection of low-dimensional spaces containing the set of residuals associated with a dictionary-based approximation. In particular, this condition can be ensured by using random projections constructed with SRHT or Gaussian matrices of sufficiently large sizes depending only logarithmically on the cardinality of the dictionary and the probability of failure.

This article has also addressed the efficient post-processing of the approximate solution from its coordinates associated with the given reduced basis (or dictionary). The extraction of the output quantity from an approximate solution can be done by a (sufficiently accurate) projection onto a low-dimensional space with a random sketching correction. This approach can be particularly important for the approximation of the linear output with an expensive extractor of the quantity of interest, as well as quadratic output and primal-dual correction.

Finally, we provided a probabilistic approach for a posteriori certification of the quality of the (random) embedding (and the associated sketch). This procedure is online-efficient and does not require operations on high dimensional vectors but only on their small sketches. It can be particularly useful for an adaptive selection of the size of a random sketching matrix since a priori bounds were revealed to be highly pessimistic.

The great applicability of the proposed methodology was realized on two benchmark problems difficult to tackle with standard methods. The experiments on the invisibility cloak benchmark proved that random sketching indeed provides high computational savings in both offline and online stages compared to the standard minres method while preserving the quality of the output. In particular, we could not even execute the classical greedy algorithm based on minres projection due to exceeding the RAM. Moreover, the improvement of numerical stability of computation (or minimization) of the residual error was also validated. It was verified experimentally that random sketching is much better suited to minres methods than to Galerkin methods. Furthermore, the procedure for a posteriori certification of the sketch was also validated in a series of experiments. It was revealed that this procedure can be used for (adaptive) selection of an (almost) optimal size of random sketching matrices in particular less by an order of magnitude than the theoretical bounds from [3].

Finally, we considered an advection-diffusion benchmark defined on a complex flow. For this problem a slow-decay of the Kolmogorov r -width was revealed, which implied the necessity to use the dictionary-based approximation. It was verified that for this problem the dictionary-based approximation provided an enhancement of the online stage in more than an order of magnitude in complexity, in about 2 times in terms of memory, and in about 4 times in terms of runtime.

Moreover, even higher computational savings could be obtained by representing the sketch of the dictionary in a more favorable format (e.g., as in [25, 29]), which we leave for the future research.

References

- [1] D. Amsallem, C. Farhat, and M. Zahr. On the robustness of residual minimization for constructing pod-based reduced-order cfd models. *21st AIAA Computational Fluid Dynamics Conference*, 2013.
- [2] D. Amsallem and B. Haasdonk. Pebl-rom: Projection-error based local reduced-order models. *Advanced Modeling and Simulation in Engineering Sciences*, 3(1):6, 2016.
- [3] O. Balabanov and A. Nouy. Randomized linear algebra for model reduction. Part I: Galerkin methods and error estimation. *arXiv preprint arXiv:1803.02602*, 2019.
- [4] P. Benner, A. Cohen, M. Ohlberger, and K. Willcox, editors. *Model Reduction and Approximation: Theory and Algorithms*. SIAM, Philadelphia, PA, 2017.
- [5] K. Bertoldi, V. Vitelli, J. Christensen, and M. van Hecke. Flexible mechanical metamaterials. *Nature Reviews Materials*, 2(11):17066, 2017.
- [6] D. A. Bistrrian and I. M. Navon. Randomized dynamic mode decomposition for nonintrusive reduced order modelling. *International Journal for Numerical Methods in Engineering*, 112(1):3–25, 2017.
- [7] S. L. Brunton, J. L. Proctor, J. H. Tu, and J. N. Kutz. Compressed sensing and dynamic mode decomposition. *Journal of Computational Dynamics*, 2:165, 2015.
- [8] A. Buhr and K. Smetana. Randomized local model order reduction. *SIAM Journal on Scientific Computing*, 40:2120–2151, 2018.
- [9] T. Bui-Thanh, K. Willcox, and O. Ghattas. Model reduction for large-scale systems with high-dimensional parametric input space. *SIAM Journal on Scientific Computing*, 30(6):3270–3288, 2008.
- [10] H. Chen and C. T. Chan. Acoustic cloaking and transformation acoustics. *Journal of Physics D: Applied Physics*, 43(11):113001, 2010.
- [11] Y. Cheng, F. Yang, J. Y. Xu, and X. J. Liu. A multilayer structured acoustic cloak with homogeneous isotropic materials. *Applied Physics Letters*, 92(15):151913, 2008.
- [12] R. A. DeVore. Nonlinear approximation and its applications. In *Multiscale, Nonlinear and Adaptive Approximation*, pages 169–201. Springer, 2009.

- [13] M. Dihlmann, S. Kaulmann, and B. Haasdonk. Online reduced basis construction procedure for model reduction of parametrized evolution systems. *IFAC Proceedings Volumes*, 45(2):112–117, 2012.
- [14] J. L. Eftang, D. J. Knezevic, and A. T. Patera. An hp certified reduced basis method for parametrized parabolic partial differential equations. *Mathematical and Computer Modelling of Dynamical Systems*, 17(4):395–422, 2011.
- [15] J. L. Eftang, A. T. Patera, and E. M. Rønquist. An ”hp” certified reduced basis method for parametrized elliptic partial differential equations. *SIAM Journal on Scientific Computing*, 32(6):3170–3200, 2010.
- [16] N. B. Erichson and C. Donovan. Randomized low-rank dynamic mode decomposition for motion detection. *Computer Vision and Image Understanding*, 146:40–50, 2016.
- [17] N. B. Erichson, L. Mathelin, S. L. Brunton, and J. N. Kutz. Randomized dynamic mode decomposition. *arXiv preprint [arXiv:1702.02912](https://arxiv.org/abs/1702.02912)*, 2017.
- [18] B. Haasdonk. Reduced basis methods for parametrized PDEs – A tutorial introduction for stationary and instationary problems. *Model reduction and approximation: theory and algorithms*, 15:65, 2017.
- [19] N. Halko, P.-G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.
- [20] C. Homescu, L. R. Petzold, and R. Serban. Error estimation for reduced-order models of dynamical systems. *SIAM Journal on Numerical Analysis*, 43(4):1693–1714, 2005.
- [21] 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. *Comptes Rendus Mathématique*, 345(8):473–478, 2007.
- [22] M. Kadic, T. Bückmann, R. Schittny, P. Gumbsch, and M. Wegener. Pentamode metamaterials with independently tailored bulk modulus and mass density. *Physical Review Applied*, 2(5):054007, 2014.
- [23] S. Kaulmann and B. Haasdonk. Online greedy reduced basis construction using dictionaries. In *VI International Conference on Adaptive Modeling and Simulation (ADMOS 2013)*, pages 365–376, 2013.
- [24] B. Kramer, P. Grover, P. Boufounos, S. Nabi, and M. Benosman. Sparse sensing and dmd-based identification of flow regimes and bifurcations in complex flows. *SIAM Journal on Applied Dynamical Systems*, 16(2):1164–1196, 2017.
- [25] L. Le Magoarou and R. Gribonval. Flexible multilayer sparse approximations of matrices and applications. *IEEE Journal of Selected Topics in Signal Processing*, 10(4):688–700, 2016.

- [26] Y. Maday and B. Stamm. Locally adaptive greedy approximations for anisotropic parameter reduced basis spaces. *SIAM Journal on Scientific Computing*, 35(6):2417–2441, 2013.
- [27] M. W. Mahoney et al. Randomized algorithms for matrices and data. *Foundations and Trends® in Machine Learning*, 3(2):123–224, 2011.
- [28] V. Rokhlin and M. Tygert. A fast randomized algorithm for overdetermined linear least-squares regression. *Proceedings of the National Academy of Sciences*, 105(36):13212–13217, 2008.
- [29] R. Rubinfeld, M. Zibulevsky, and M. Elad. Double sparsity: Learning sparse dictionaries for sparse signal approximation. *IEEE Transactions on signal processing*, 58(3):1553–1564, 2009.
- [30] K. Smetana, O. Zahm, and A. T. Patera. Randomized residual-based error estimators for parametrized equations. *arXiv preprint [arXiv:1807.10489](https://arxiv.org/abs/1807.10489)*, 2018.
- [31] V. N. Temlyakov. Nonlinear Kolmogorov widths. *Mathematical Notes*, 63(6):785–795, 1998.
- [32] J. A. Tropp and A. C. Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Transactions on information theory*, 53(12):4655–4666, 2007.
- [33] D. P. Woodruff et al. Sketching as a tool for numerical linear algebra. *Foundations and Trends® in Theoretical Computer Science*, 10(1–2):1–157, 2014.
- [34] O. Zahm and A. Nouy. Interpolation of inverse operators for preconditioning parameter-dependent equations. *SIAM Journal on Scientific Computing*, 38(2):1044–1074, 2016.

Appendix

Here we list the proofs of propositions from the paper.

Proof of Proposition 3.1. The statement of the proposition follows directly from the definitions of the constants $\zeta_r(\mu)$ and $\iota_r(\mu)$, that imply

$$\zeta_r(\mu)\|\mathbf{u}(\mu) - \mathbf{u}_r(\mu)\|_U \leq \|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'} \leq \|\mathbf{r}(\mathbf{P}_{U_r}\mathbf{u}(\mu); \mu)\|_{U'} \leq \iota_r(\mu)\|\mathbf{u}(\mu) - \mathbf{P}_{U_r}\mathbf{u}(\mu)\|_U. \quad \square$$

Proof of Proposition 3.2. The proof follows the one of Proposition 3.1. □

Proof of Proposition 3.3. By the assumption on Θ , we have that

$$\sqrt{1 - \varepsilon}\|\mathbf{A}(\mu)\mathbf{x}\|_{U'} \leq \|\mathbf{A}(\mu)\mathbf{x}\|_{U'}^{\Theta} \leq \sqrt{1 + \varepsilon}\|\mathbf{A}(\mu)\mathbf{x}\|_{U'}$$

holds for all $\mathbf{x} \in \text{span}\{\mathbf{u}(\mu)\} + U_r$. Then (19) follows immediately. □

Proof of Proposition 3.5. Let $\mathbf{a} \in \mathbb{K}^r$ and $\mathbf{x} := \mathbf{U}_r\mathbf{a}$. Then

$$\frac{\|\mathbf{V}_r^{\Theta}(\mu)\mathbf{a}\|}{\|\mathbf{a}\|} = \frac{\|\Theta\mathbf{R}_U^{-1}\mathbf{A}(\mu)\mathbf{U}_r\mathbf{a}\|}{\|\mathbf{a}\|} = \frac{\|\mathbf{A}(\mu)\mathbf{x}\|_{U'}^{\Theta}}{\|\mathbf{x}\|_U^{\Theta}}.$$

Since Θ is an ε -embedding for U_r , we have

$$\sqrt{1 - \varepsilon}\|\mathbf{x}\|_U \leq \|\mathbf{x}\|_U^{\Theta} \leq \sqrt{1 + \varepsilon}\|\mathbf{x}\|_U.$$

The statement of the proposition follows immediately. □

Proof of Proposition 4.1. Define

$$\mathcal{D}_K^{(i)} = \arg \min_{\#\mathcal{D}_K=K_i} \sup_{\mathbf{u} \in \mathcal{M}^{(i)}} \min_{W_{r_i} \in \mathcal{L}_{r_i}(\mathcal{D}_K)} \|\mathbf{u} - \mathbf{P}_{W_{r_i}}\mathbf{u}\|_U,$$

and

$$\mathcal{D}_K^* = \bigcup_{i=1}^l \mathcal{D}_K^{(i)}.$$

The following relations hold:

$$\begin{aligned} \sum_{i=1}^l \sigma_{r_i}(\mathcal{M}^{(i)}; K_i) &= \sum_{i=1}^l \sup_{\mathbf{u} \in \mathcal{M}^{(i)}} \min_{W_{r_i} \in \mathcal{L}_{r_i}(\mathcal{D}_K^{(i)})} \|\mathbf{u} - \mathbf{P}_{W_{r_i}}\mathbf{u}\|_U \\ &\geq \sup_{\mu \in \mathcal{P}} \sum_{i=1}^l \min_{W_{r_i} \in \mathcal{L}_{r_i}(\mathcal{D}_K^{(i)})} \|\mathbf{u}^{(i)}(\mu) - \mathbf{P}_{W_{r_i}}\mathbf{u}^{(i)}(\mu)\|_U \\ &\geq \sup_{\mu \in \mathcal{P}} \min_{W_r \in \mathcal{L}_r(\mathcal{D}_K^*)} \sum_{i=1}^l \|\mathbf{u}^{(i)}(\mu) - \mathbf{P}_{W_r}\mathbf{u}^{(i)}(\mu)\|_U \\ &\geq \sup_{\mu \in \mathcal{P}} \min_{W_r \in \mathcal{L}_r(\mathcal{D}_K^*)} \left\| \sum_{i=1}^l \mathbf{u}^{(i)}(\mu) - \mathbf{P}_{W_r} \sum_{i=1}^l \mathbf{u}^{(i)}(\mu) \right\|_U \geq \sigma_r(\mathcal{M}; K). \end{aligned}$$

□

Proof of Proposition 4.3. Let $U_r^*(\mu) := \arg \min_{W_r \in \mathcal{L}_r(\mathcal{D}_K)} \|\mathbf{u}(\mu) - \mathbf{P}_{W_r} \mathbf{u}(\mu)\|_U$. By definition of $\mathbf{u}_r(\mu)$ and constants $\zeta_{r,K}(\mu)$ and $\iota_{r,K}(\mu)$,

$$\begin{aligned} \zeta_{r,K}(\mu) \|\mathbf{u}(\mu) - \mathbf{u}_r(\mu)\|_U &\leq \|\mathbf{r}(\mathbf{u}_r(\mu); \mu)\|_{U'} \leq \|\mathbf{r}(\mathbf{P}_{U_r^*(\mu)} \mathbf{u}(\mu); \mu)\|_{U'} \\ &\leq \iota_{r,K}(\mu) \|\mathbf{u}(\mu) - \mathbf{P}_{U_r^*(\mu)} \mathbf{u}(\mu)\|_U, \end{aligned}$$

which ends the proof. \square

Proof of Proposition 4.4. The proof exactly follows the one of Proposition 4.3 by replacing $\|\cdot\|_{U'}$ with $\|\cdot\|_{U'}^\ominus$. \square

Proof of Proposition 4.5. We have that

$$\sqrt{1-\varepsilon} \|\mathbf{A}(\mu)\mathbf{x}\|_{U'} \leq \|\mathbf{A}(\mu)\mathbf{x}\|_{U'}^\ominus \leq \sqrt{1+\varepsilon} \|\mathbf{A}(\mu)\mathbf{x}\|_{U'}$$

holds for all $\mathbf{x} \in \text{span}\{\mathbf{u}(\mu)\} + W_r$ with $W_r \in \mathcal{L}_r(\mathcal{D}_K)$. The statement of the proposition then follows directly from the definitions of $\zeta_{r,K}(\mu)$, $\iota_{r,K}(\mu)$, $\zeta_{r,K}^\ominus(\mu)$ and $\iota_{r,K}^\ominus(\mu)$. \square

Proof of Proposition 4.6. Let $\mathbf{U}_r(\mu) \in \mathbb{K}^{n \times r}$ be a matrix whose column vectors are selected from the dictionary \mathcal{D}_K . Let $\mathbf{x} \in \mathbb{K}^r$ be an arbitrary vector and $\mathbf{w}(\mu) := \mathbf{U}_r(\mu)\mathbf{x}$. Let $\mathbf{z}(\mu) \in \mathbb{K}^K$ with $\|\mathbf{z}(\mu)\|_0 \leq r$ be a sparse vector such that $\mathbf{U}_K \mathbf{z}(\mu) = \mathbf{w}(\mu)$. Then

$$\begin{aligned} \frac{\|\mathbf{V}_r^\ominus(\mu)\mathbf{x}\|}{\|\mathbf{x}\|} &= \frac{\|\Theta \mathbf{R}_U^{-1} \mathbf{A}(\mu) \mathbf{U}_r(\mu) \mathbf{x}\|}{\|\mathbf{x}\|} = \frac{\|\mathbf{A}(\mu) \mathbf{w}(\mu)\|_{U'}^\ominus}{\|\mathbf{x}\|} = \frac{\|\mathbf{A}(\mu) \mathbf{w}(\mu)\|_{U'}^\ominus}{\|\mathbf{w}(\mu)\|_U} \frac{\|\mathbf{w}(\mu)\|_U}{\|\mathbf{x}\|} \\ &\geq \zeta_{r,K}^\ominus(\mu) \frac{\|\mathbf{U}_r(\mu)\mathbf{x}\|_U}{\|\mathbf{x}\|} = \zeta_{r,K}^\ominus(\mu) \frac{\|\mathbf{U}_K \mathbf{z}(\mu)\|_U}{\|\mathbf{z}(\mu)\|} \geq \zeta_{r,K}^\ominus(\mu) \Sigma_{r,K}^{\min}. \end{aligned}$$

Similarly,

$$\frac{\|\mathbf{V}_r^\ominus(\mu)\mathbf{x}\|}{\|\mathbf{x}\|} \leq \iota_{r,K}^\ominus(\mu) \frac{\|\mathbf{U}_r(\mu)\mathbf{x}\|_U}{\|\mathbf{x}\|} \leq \iota_{r,K}^\ominus(\mu) \Sigma_{r,K}^{\max}.$$

The statement of the proposition follows immediately. \square

Proof of Proposition 5.3. Denote $\mathbf{x} := \mathbf{R}_U^{-1} \mathbf{1}(\mu) / \|\mathbf{1}(\mu)\|_{U'}$ and $\mathbf{y} := (\mathbf{u}_r(\mu) - \mathbf{w}_p(\mu)) / \|\mathbf{u}_r(\mu) - \mathbf{w}_p(\mu)\|_U$. Let us consider $\mathbb{K} = \mathbb{C}$, which also accounts for the real case, $\mathbb{K} = \mathbb{R}$. Let

$$\omega := \frac{\langle \mathbf{x}, \mathbf{y} \rangle_U - \langle \mathbf{x}, \mathbf{y} \rangle_U^\ominus}{|\langle \mathbf{x}, \mathbf{y} \rangle_U - \langle \mathbf{x}, \mathbf{y} \rangle_U^\ominus|}.$$

Observe that $|\omega| = 1$ and $\langle \mathbf{x}, \omega \mathbf{y} \rangle_U - \langle \mathbf{x}, \omega \mathbf{y} \rangle_U^\ominus$ is a real number.

By a union bound for the probability of success, Θ is an ε -embedding for $\text{span}(\mathbf{x} + \omega\mathbf{y})$ and $\text{span}(\mathbf{x} - \omega\mathbf{y})$ with probability at least $1 - 2\delta$. Then, using the parallelogram identity we obtain

$$\begin{aligned}
4|\langle \mathbf{x}, \mathbf{y} \rangle_U - \langle \mathbf{x}, \mathbf{y} \rangle_U^\Theta| &= |4\langle \mathbf{x}, \omega\mathbf{y} \rangle_U - 4\langle \mathbf{x}, \omega\mathbf{y} \rangle_U^\Theta| \\
&= \left| \|\mathbf{x} + \omega\mathbf{y}\|_U^2 - \|\mathbf{x} - \omega\mathbf{y}\|_U^2 + 4\text{Im}(\langle \mathbf{x}, \omega\mathbf{y} \rangle_U) \right. \\
&\quad \left. - \left(\|\mathbf{x} + \omega\mathbf{y}\|_U^\Theta{}^2 - \|\mathbf{x} - \omega\mathbf{y}\|_U^\Theta{}^2 + 4\text{Im}(\langle \mathbf{x}, \omega\mathbf{y} \rangle_U^\Theta) \right) \right| \\
&= \left| \|\mathbf{x} + \omega\mathbf{y}\|_U^2 - (\|\mathbf{x} + \omega\mathbf{y}\|_U^\Theta)^2 - (\|\mathbf{x} - \omega\mathbf{y}\|_U^2 - (\|\mathbf{x} - \omega\mathbf{y}\|_U^\Theta)^2) \right. \\
&\quad \left. - 4\text{Im}(\langle \mathbf{x}, \omega\mathbf{y} \rangle_U - \langle \mathbf{x}, \omega\mathbf{y} \rangle_U^\Theta) \right| \\
&\leq \varepsilon\|\mathbf{x} + \omega\mathbf{y}\|_U^2 + \varepsilon\|\mathbf{x} - \omega\mathbf{y}\|_U^2 = 4\varepsilon.
\end{aligned}$$

We conclude that relation (52) holds with probability at least $1 - 2\delta$. \square

Proof of Proposition 5.4. We can use a similar proof as in Proposition 5.3 with the fact that if Θ is an ε -embedding for $L + U_r + W_p$, then it satisfies the ε -embedding property for $\text{span}(\mathbf{x} + \omega\mathbf{y})$ and $\text{span}(\mathbf{x} - \omega\mathbf{y})$. \square

Proof of Proposition 6.2. Using Proposition 5.3 with $\mathbf{l}(\mu) := \mathbf{R}_U\mathbf{x}$, $\mathbf{u}_r(\mu) := \mathbf{y}$, $\mathbf{w}_p(\mu) := \mathbf{0}$, $\Theta := \Theta^*$, $\varepsilon := \varepsilon^*$ and $\delta := \delta^*$, we have

$$\mathbb{P}(|\langle \mathbf{x}, \mathbf{y} \rangle_U - \langle \mathbf{x}, \mathbf{y} \rangle_U^{\Theta^*}| \leq \varepsilon^*\|\mathbf{x}\|_U\|\mathbf{y}\|_U) \geq 1 - 2\delta^*, \quad (67)$$

from which we deduce that

$$\begin{aligned}
|\langle \mathbf{x}, \mathbf{y} \rangle_U^{\Theta^*} - \langle \mathbf{x}, \mathbf{y} \rangle_U^\Theta| - \varepsilon^*\|\mathbf{x}\|_U\|\mathbf{y}\|_U &\leq |\langle \mathbf{x}, \mathbf{y} \rangle_U - \langle \mathbf{x}, \mathbf{y} \rangle_U^\Theta| \\
&\leq |\langle \mathbf{x}, \mathbf{y} \rangle_U^{\Theta^*} - \langle \mathbf{x}, \mathbf{y} \rangle_U^\Theta| + \varepsilon^*\|\mathbf{x}\|_U\|\mathbf{y}\|_U
\end{aligned} \quad (68)$$

holds with probability at least $1 - 2\delta^*$. In addition,

$$\mathbb{P}(|\|\mathbf{x}\|_U^2 - (\|\mathbf{x}\|_U^\Theta)^2| \leq \varepsilon^*\|\mathbf{x}\|_U^2) \geq 1 - \delta^* \quad (69)$$

and

$$\mathbb{P}(|\|\mathbf{y}\|_U^2 - (\|\mathbf{y}\|_U^\Theta)^2| \leq \varepsilon^*\|\mathbf{y}\|_U^2) \geq 1 - \delta^*. \quad (70)$$

The statement of the proposition can be now derived by combining (68) to (70) and using a union bound argument. \square

Proof of Proposition 6.3. Observe that

$$\omega = \max \left\{ 1 - \min_{\mathbf{x} \in V/\{\mathbf{0}\}} \left(\frac{\|\mathbf{x}\|_U^\Theta}{\|\mathbf{x}\|_U} \right)^2, \max_{\mathbf{x} \in V/\{\mathbf{0}\}} \left(\frac{\|\mathbf{x}\|_U^\Theta}{\|\mathbf{x}\|_U} \right)^2 - 1 \right\}.$$

Let us make the following assumption:

$$1 - \min_{\mathbf{x} \in V/\{\mathbf{0}\}} \left(\frac{\|\mathbf{x}\|_U^\Theta}{\|\mathbf{x}\|_U} \right)^2 \geq \max_{\mathbf{x} \in V/\{\mathbf{0}\}} \left(\frac{\|\mathbf{x}\|_U^\Theta}{\|\mathbf{x}\|_U} \right)^2 - 1.$$

For the alternative case the proof is similar.

Next, we show that $\bar{\omega}$ is an upper bound for ω with probability at least $1 - \delta^*$. Define $\mathbf{x}^* := \arg \min_{\mathbf{x} \in V/\{\mathbf{0}\}, \|\mathbf{x}\|_U=1} \|\mathbf{x}\|_U^{\Theta}$. By definition of Θ^* ,

$$1 - \varepsilon^* \leq (\|\mathbf{x}^*\|_U^{\Theta^*})^2 \quad (71)$$

holds with probability at least $1 - \delta^*$. If (71) is satisfied, we have

$$\bar{\omega} \geq 1 - (1 - \varepsilon^*) \min_{\mathbf{x} \in V/\{\mathbf{0}\}} \left(\frac{\|\mathbf{x}\|_U^{\Theta}}{\|\mathbf{x}\|_U^{\Theta^*}} \right)^2 \geq 1 - (1 - \varepsilon^*) \left(\frac{\|\mathbf{x}^*\|_U^{\Theta}}{\|\mathbf{x}^*\|_U^{\Theta^*}} \right)^2 \geq 1 - (\|\mathbf{x}^*\|_U^{\Theta})^2 = \omega.$$

□

Proof of Proposition 6.4. By definition of ω and the assumption on Θ^* , for all $\mathbf{x} \in V$, it holds

$$|\|\mathbf{x}\|_U^2 - (\|\mathbf{x}\|_U^{\Theta^*})^2| \leq \omega^* \|\mathbf{x}\|_U^2, \quad \text{and} \quad |\|\mathbf{x}\|_U^2 - (\|\mathbf{x}\|_U^{\Theta})^2| \leq \omega \|\mathbf{x}\|_U^2.$$

The above relations and the definition (57) of $\bar{\omega}$ yield

$$\bar{\omega} \leq \max \left\{ 1 - (1 - \varepsilon^*) \frac{1 - \omega}{1 + \omega^*}, (1 + \varepsilon^*) \frac{1 + \omega}{1 - \omega^*} - 1 \right\} = (1 + \varepsilon^*) \frac{1 + \omega}{1 - \omega^*} - 1,$$

which ends the proof.

□