Numerical Linear Algebra with Applications

WILEY

RESEARCH ARTICLE OPEN ACCESS

Stabilized Krylov Subspace Recurrences via Randomized Sketching

Valeria Simoncini[1,2] | YiHong Wang[3]

[1]Dipartimento di Matematica, (AM)[2], Alma Mater Studiorum — Università di Bologna, Bologna, Italy | [2]IMATI-CNR, Pavia, Italy | [3]School of Mathematical Sciences, Nanjing Normal University, Nanjing, P.R. China

Correspondence: Valeria Simoncini (valeria.simoncini@unibo.it)

Received: 6 December 2024 | **Revised:** 13 March 2025 | **Accepted:** 5 May 2025

Funding: This work was supported by the National Natural Science Foundation of China (Grant No. 12131004), the INdAM and the MUR. This work was carried out during a one-year visit of Y.W. at the University of Bologna, granted by a China Scholarship Council (File No. 202306860051).

Keywords: krylov decomposition | lyapunov equation | matrix function evaluations | polynomial krylov subspaces | sketching

ABSTRACT

Recurrences building orthonormal bases for polynomial Krylov spaces have been classically used for approximation purposes in various numerical linear algebra contexts. Variants aiming to limit memory and computational costs by using truncated recurrences often have convergence constraints. Recently, randomized linear algebra strategies have been devised that drastically improve the performance of these variants, while keeping the costs low. We provide a unifying framework for analyzing a large class of Krylov subspace methods, including randomization-enhanced strategies, based on Krylov decompositions. This framework allows us to identify the key quantities — the canonical angles among the Krylov subspace basis vectors — for assessing the effectiveness of the randomized strategy. Moreover, it also allows us to analyze the spectral properties of the projected problem. Our results are illustrated with experiments using the nonsymmetric Lanczos iteration, which is an inherently three-term recurrence, so that no truncation needs to be performed. Hence, randomized procedures applied to the Lanczos method may be viewed as a way to stabilize the approximation.

MSC2020 Classification: 68W20, 65F10, 65F45, 65F60

1 | Introduction

Projection methods based on polynomial Krylov subspaces for solving numerical linear algebra problems have been largely explored in the past few decades. These are characterized by recurrences, with either few or many terms, that explicitly build basis vectors for the sought-after space,

$$\mathcal{K}_k(A, c) := \text{span}\{c, Ac, \ldots, A^{k-1}c\} \tag{1.1}$$

where $A \in \mathbb{R}^{n\times n}$ and $c \in \mathbb{R}^n$, $c \neq 0$. This *Krylov* space $\mathcal{K}_k(A, c)$ is used as the approximation space in which an approximate solution to the given problem is sought. The latter is usually obtained by projecting the original problem, so as to obtain a smaller problem of the same type, see, for example, [1, 2].

For large n and general nonsymmetric A, constructing an orthogonal basis for $\mathcal{K}_k(A, c)$ can be very expensive both in terms of memory and computational costs. A large number of strategies have been proposed to overcome this problem, among which are restarting, truncating, and different forms of optimality other than orthogonality, see, for example, [3–8] and their references. Some of these appealing strategies, such as truncation in the number of orthonormal vectors, may suffer from a lack of good

approximation quality in case truncation is too severe. In other words, if a basis is constructed that is only locally orthogonal, this basis may be too ill conditioned to be of any use; see, for example, [9]. As opposed to the methods just discussed, the Lanczos method generates a Krylov subspace basis using a three-term recurrence. This is made possible by imposing a biorthogonality constraint of the generated basis vectors with respect to the basis of a Krylov subspace constructed with A^T (see, e.g., [10, sec. 10.5.5], [11]). It is well known that biorthogonality is quickly lost in practical computations, and that each of the two bases can lose linear independence. Hence, Lanczos iterations are akin to a truncated recurrence, in which biorthogonality is gradually lost.

Very recently, randomized techniques for Krylov subspaces have been proposed to significantly contribute to mitigating the side effects of lack of full orthogonality; they go under the name of *sketching*. In these strategies, orthogonality is imposed not only locally among subsequent basis vectors, but also among a sample of rows of all given basis vectors. In its simplest form, this is carried out by randomly selecting a subsample of some combination of the rows of all basis vectors, which will form a reduced orthonormal basis. This way, it has been experimentally observed that as the iterations proceed, the generated space keeps growing, and the approximation can progress. The idea has been very successful in many different problems, such as algebraic linear systems [12–14], eigenvalue problems [15], least squares problems [16], matrix function evaluations and model order reduction [17–19], matrix equations [20]; we also refer to the references cited in these articles for a broader recollection. This paper aims to contribute to the understanding and formalization of this procedure within the Krylov subspace context, adding to the overall less rich literature in this direction.

We propose a unifying formulation of Krylov subspace recurrences, based on the concept of *Krylov decomposition*, initially introduced by Pete Stewart in Reference [21]. Many of the known recurrences, such as truncated, restarted, biorthogonal, etc. can be formalized as *equivalent* Krylov decompositions with proper translations and similarity transformations. Our contribution is manifold. First, we prove that sketched Krylov subspaces as well satisfy equivalent Krylov decompositions. By using this new result, we derive a linear independence property of the sketched basis that ensures the growth of the approximation space; in turn, this allows the progress of the underlying approximation procedure. We also deduce that under certain hypotheses, sketching may create spurious vectors that fall outside the ideal Krylov subspace, thus, possibly failing in correcting misconvergence of the truncated method. Third, we prove that the eigenvalues of the sketched projected problem remain in a region close to that characterizing the spectral properties of the reduced matrix of the fully orthogonal basis. All these results are displayed by using the nonsymmetric Lanczos iteration, in which sketching is performed only on the recurrence that forms the Krylov decomposition for the matrix A. The Lanczos approach is inherently three-term, so that no truncation parameter needs to be tuned. The benefit of this feature is illustrated in our numerical experiments. In particular, not having to cure possible misconvergence due to truncation, the sketching procedure may be viewed as a stabilization procedure, obtained by injecting more linear independence in the constructed basis. As samples of typical timely problems on which we test the potential of the sketched Lanczos method, we consider the approximation of matrix function evaluations, and the approximate solution of the Lyapunov equation.

Finally, we would like to stress that our results assume exact precision arithmetic. Nonetheless, our discussion also addresses behaviors in finite precision arithmetic. For instance, the sketching strategy appears to be particularly effective in mitigating the effects of loss of biorthogonality caused by round-off errors in the nonsymmetric Lanczos recurrence. Indeed, in finite precision arithmetic k iterations of the Lanczos method may be viewed as an exact precision arithmetic recurrence with k terms (building a non-biorthogonal basis), to which our analysis can be applied. We also notice that in finite precision arithmetic, the basis of a Krylov decomposition may lose linear independence, and in that case, our results no longer hold. In particular, in finite precision arithmetic, spurious vectors may be generated whenever an equivalent Krylov decomposition is obtained after sketching a matrix that is not numerically full rank.

An outline of the paper is as follows. In Section 2 we recall the Krylov decomposition and its properties, indicating some of the major Krylov-based procedures it encompasses. In Section 3, we survey some of the key Krylov decomposition methods that will be used in later sections, while in Section 4, we derive some general spectral properties associated with the lack of orthogonality of the basis. In Section 5 we recall the randomized strategies (sketching) for Krylov bases, while in Section 6 we prove that the sketched basis also determines a Krylov decomposition, with a well-conditioned basis, as long as the original basis is full rank. In Section 7 we show that sketched bases also carry good properties in terms of the spectral eigenvalue location, compared to generic Krylov decompositions. Finally, Sections 8 and 9 collect a number of computational experiments to illustrate our theoretical results: We focus on two timely problems: Matrix functions evaluation applied to a vector, and the solution of Lyapunov matrix equations, respectively.

Notation. The Euclidean norm for vectors and the corresponding induced norm for matrices is used. Here and in the following, x^* denotes the conjugate transpose of a complex vector x, while x^T denotes the real transpose. With I we denote the identity matrix, whose dimension is clear from the context, and with e_j its jth column. We denote with spec(A) the set of the eigenvalues of A. The field of values of an $n \times n$ matrix A is given by

$$\mathcal{W}(A) = \{\zeta \in \mathbb{C} : \ \zeta = (x^*Ax)/(x^*x), \ 0 \neq x \in \mathbb{C}^n\}$$

We use the notation $\mathcal{O}(\tau)$ for quantities that are of the same order of magnitude as τ.

2 | Krylov Decompositions

We first recall the general definition of Krylov decompositions, introduced in Reference [21]. Then we highlight how this class of decompositions can be exploited in our context.

Definition 2.1 ([21]). A Krylov decomposition of order k is a relation of the form

$$AU_k = U_k B_k + u_{k+1} b_{k+1}^* \tag{2.1}$$

where B_k is of order k and is called Rayleigh quotient of the decomposition, b_{k+1} is a vector with k components, and the $k+1$ columns of $[U_k, u_{k+1}]$ are linearly independent.

We first collect some properties associated with (2.1), and later on, we will show that these properties are shared by several Krylov subspace methods that are currently used, including certain randomization-based strategies. It is important to realize that the linear independence hypothesis is far less stringent than orthogonality. In exact arithmetic many of the known Krylov-based recurrences do satisfy this latter condition, although the matrix $[U_k, u_{k+1}]$ may become arbitrarily ill conditioned in finite precision arithmetic.

As discussed in Reference [21], Krylov decompositions are closed under translation and similarity transformations. Translations allow us to change the vector u_{k+1}, while similarity transformations act on B_k, U_k and b_{k+1}. Translations act as follows. Let

$$\eta_k \hat{u}_{k+1} := u_{k+1} - U_k g_k, \qquad \eta_k \neq 0$$

note that $\hat{u}_{k+1}$ is ensured to always contain a nonzero component onto the extra vector u_{k+1}. Substituting u_{k+1} in Equation (2.1) we obtain

$$AU_k = U_k(B_k + g_k b_{k+1}^*) + \hat{u}_{k+1}\eta_k b_{k+1}^*$$

Hence, translations perform a rank-one modification of the Rayleigh quotient matrix.

Given $R \in \mathbb{R}^{k\times k}$ nonsingular, similarity transformations can be written as

$$AU_k R^{-1} = U_k R^{-1}(RB_k R^{-1}) + u_{k+1}(b_{k+1}^* R^{-1})$$

so that

$$A\hat{U}_k = \hat{U}_k \hat{B}_k + u_{k+1}\hat{b}_{k+1}^*$$

with

$$\hat{U}_k = U_k R^{-1}, \quad \hat{B}_k = RB_k R^{-1}, \quad \hat{b}_{k+1}^* = b_{k+1}^* R^{-1}$$

Stewart in Reference [21] calls *equivalent* those Krylov decompositions that can be derived from each other by a sequence of translations and similarities. Collecting both transformations, and also setting $\hat{g}_k = Rg_k$, we can write

$$A\hat{U}_k = \hat{U}_k(\hat{B}_k + \hat{g}_k\hat{b}_{k+1}^*) + \hat{u}_{k+1}\eta_k\hat{b}_{k+1}^* \tag{2.2}$$

which is again a Krylov decomposition. This form will be used throughout this paper.

For a general nonsymmetric A, the Arnoldi relation [2] is a special case of Krylov decomposition where the matrix $[U_k, u_{k+1}]$ has orthonormal columns, B_k is upper Hessenberg and $b_{k+1} = \beta_k e_k^T$. Many other recurrences can be brought to a Krylov decomposition form; some of them will be discussed later in greater detail. For instance,

1. The truncated Arnoldi recurrence generates the relation (2.1) with $[U_k, u_{k+1}]$ only having local orthogonality, and B_k is upper Hessenberg and banded [2, section 6.4.2];
2. The restarted Arnoldi recurrence generates (2.1) with $[U_k, u_{k+1}]$ having blocks of orthonormal columns (with block size corresponding to the restarting parameter), and B_k is almost block diagonal [22];
3. Recurrences using polynomials of degree i in the form $\phi_i(A)c$, that generate bases $\{\phi_0(A)c, \phi_1(A)c, \ldots, \phi_{k-1}(A)c\}$. These are obtained, for instance, by Newton or Chebyshev polynomial sequences, see, for example, [23] and references therein;
4. The nonsymmetric Lanczos algorithm determines two Krylov subspaces, for A and A^T, respectively, each having the form (2.1).

Other recent strategies include communication-avoiding recurrences [24] and limited-memory methods for symmetric matrices in procedures where the basis is anyway required [25]. In addition, we will show that sketched truncated Arnoldi methods can also be completely characterized as an *equivalent* Krylov decomposition. All these methods use b_{k+1} as a multiple of the canonical vector e_k, and this plays a relevant role in our analysis.

We aim to deepen our understanding of the role of the Rayleigh quotient matrix B_k and of its translated version, whenever the columns of U_{k+1} are not orthonormal, but still linearly independent. To this end, we recall the definition of (1)-inverse of M, which is a matrix satisfying the first of the four conditions of the Moore-Penrose pseudo-inverse, namely $MM^+M = M$. The (1)-inverse is a family of parametrized matrices [26, Th. 6.3.3], characterized by being equation-solving matrices, thus usually also called *equation solving generalized inverses* [26, Th. 6.2.2]. This equation-solving property will be used in our analysis to explicitly determine a representation of the reduced matrix in terms of the quantities used in the Krylov decomposition. If M is full column rank, a member of the family is given by $M^+ = (M^TM)^{-1}M^T$, which is the Moore-Penrose generalized inverse. We will use this form in the following lemma, whereas other instances of the (1)-inverse will be used in Sections 3.1 and 3.2. For later convenience, in the next lemma, we explicitly write $[U_k, u_{k+1}]^+$ in terms of its partitioned blocks; see also [26, Th. 6.3.4].

Lemma 2.2. *Assume that $U_{k+1} = [U_k, u_{k+1}]$ is full column rank and let Π_k be the projection operator $\Pi_k = I - u_{k+1}(u_{k+1}^Tu_{k+1})^{-1}u_{k+1}^T$. Then the matrix $U_k^T\Pi_kU_k$ is nonsingular, and a (1)-inverse U_{k+1}^+ can be written as $U_{k+1}^+ = [X_k^T; y_k^T]$ with*

$$X_k^T = (U_k^T\Pi_kU_k)^{-1}U_k^T\Pi_k \quad \text{and}$$
$$y_k^T = u_{k+1}^T(I - U_k(U_k^T\Pi_kU_k)^{-1}U_k^T\Pi_k)/(u_{k+1}^Tu_{k+1})$$

Proof. It holds that $U_{k+1}^+ = (U_{k+1}^TU_{k+1})^{-1}U_{k+1}^T$. Using $U_{k+1}^+ = [X_k^T; y_k^T]$ we can write $([U_k, u_{k+1}]^T[U_k, u_{k+1}])[U_k, u_{k+1}]^+ = [U_k, u_{k+1}]^T$, that is

$$\begin{bmatrix} U_k^TU_k & U_k^Tu_{k+1} \\ u_{k+1}^TU_k & u_{k+1}^Tu_{k+1} \end{bmatrix} \begin{bmatrix} X_k^T \\ y_k^T \end{bmatrix} = \begin{bmatrix} U_k^T \\ u_{k+1}^T \end{bmatrix}$$

Obtaining y_k^T and substituting, we get $U_k^T \Pi U_k X_k^T = U_k^T \Pi$ with $U_k^T \Pi U_k$ nonsingular. Indeed, we first notice that $U_k^T \Pi U_k = (U_k^T \Pi)(\Pi U_k)$. Moreover, if $\Pi U_k d = 0$ for some d, then it must be $[U_k, u_{k+1}][d; -(u_{k+1}^T U_k d)/(u_{k+1}^T u_{k+1})] = 0$. From the full rank assumption, we obtain that $d = 0$. The final result thus follows. □

With U_k as in Equation (2.1), we have that $X_k^T u_{k+1} = 0$ with X_k defined in Lemma 2.2. Hence, multiplying (2.1) from the left by X_k^T we obtain

$$B_k = X_k^T A U_k = (U_k^T \Pi_k U_k)^{-1} U_k^T \Pi_k A U_k \tag{2.3}$$

This general relation shows that B_k results from an oblique projection of A, that implicitly orthogonalizes U_k with respect to u_{k+1}; see, [2, section 6.4.2] for a similar result for truncated bases. We also stress that the only hypothesis employed is that (2.1) is a Krylov decomposition, that is $[U_k, u_{k+1}]$ is full column rank.

In the following sections, we will show that several classical and recent techniques can be formulated within this framework, for specifically selected transformations.

3 | Some Recurrences Generating a Krylov Subspace

In this section, we recall several recurrences that generate a Krylov subspace and that fit the Krylov decomposition framework.

Given A and c, the classical Arnoldi iteration generates a sequence of mutually orthonormal vectors $\{\boldsymbol{u}_1, \boldsymbol{u}_2, \dots\}$ with $\boldsymbol{u}_1 = c/\|c\|$ by means of the following iteration, for $k = 1, 2, \dots,$

$$u_{k+1} = A\boldsymbol{u}_k \tag{3.1}$$

$$\text{for } \; j = 1, \dots, k, \quad u_{k+1} = u_{k+1} - \boldsymbol{u}_j(\boldsymbol{u}_j^T u_{k+1}), \quad \text{end} \tag{3.2}$$

$$\boldsymbol{u}_{k+1} := u_{k+1}/\|u_{k+1}\| \tag{3.3}$$

The orthogonalization coefficients are stored in the matrix H_k with entries $h_{j,k} = \boldsymbol{u}_j^T u_{k+1}$, and $h_{k+1,k} = \|u_{k+1}\|$. This form corresponds to a Krylov decomposition: After k iterations the matrix $\boldsymbol{U}_{k+1} = [\boldsymbol{u}_1, \dots, \boldsymbol{u}_{k+1}]$ satisfies the so-called Arnoldi relation[1]

$$A\boldsymbol{U}_k = \boldsymbol{U}_k H_k + \boldsymbol{u}_{k+1} h_{k+1,k} e_k^T$$

with $[\boldsymbol{U}_k, \boldsymbol{u}_{k+1}]$ having orthonormal columns and H_k having upper Hessenberg structure. From this recurrence, it follows that $\text{Range}(\boldsymbol{U}_k) = \mathcal{K}_k(A, c)$ and $\text{Range}(\boldsymbol{U}_k) \subseteq \text{Range}(\boldsymbol{U}_{k+1})$.

As the dimension of the subspace $\mathcal{K}_k(A, c)$ increases, the orthogonalization step will dramatically increase the computational cost. This has led to the development of several memory-saving strategies, some of which are discussed below.

3.1 | Truncated and Restarted Arnoldi Iterations

One way to tackle memory constraints is to truncate the Arnoldi recurrence, that is, orthogonalize the next $\widehat{u}_{k+1}$ only with respect to the previous k_t basis vectors, where k_t is some fixed value less than k, selected a priori. This method yields the *truncated* Arnoldi relation [2, section 6.4.2]

$$AU_k = U_k H_k + u_{k+1} h_{k+1,k} e_k^T \tag{3.4}$$

where U_k is only locally orthogonal, that is, $u_i^T u_j = 0$ for $|i - j| < k_t$ and $i \neq j$. Hence, H_k is upper Hessenberg and *upper banded*, with bandwidth equal to k_t.

This relation is precisely the Krylov decomposition introduced in Section 2, for specific choices of B_k and b_{k+1}. For any full column rank matrix $Z_k \in \mathbb{R}^{n \times k}$ such that $Z_k^T u_{k+1} = 0$, it holds that $H_k = (Z_k^T U_k)^{-1} Z_k^T A U_k$. In [2, Prop. 6.8] it was suggested to select Z_k as $Z_k = (I - u_{k+1} u_{k+1}^T) U_k$. Here Z_k corresponds to the (1)-inverse block associated with B_k in Equation (2.3), and it can be similarly shown that $Z_k^T U_k = Z_k^T Z_k$ and that it is nonsingular.

A more structured truncation is obtained by restarting (cycling) the fully orthogonal Arnoldi procedure after k steps with the next basis vector $\boldsymbol{u}_{k+1}$, so that in later iterations orthogonality with respect to all previous vectors is relaxed. This procedure is what the system solver FOM does at restart time. If one were to collect two consecutive k-dimensional cycles, then one would obtain the following "truncated" Arnoldi matrix relation

$$A[\boldsymbol{U}_k^{(0)}, -\boldsymbol{U}_k^{(1)}] = [\boldsymbol{U}_k^{(0)}, -\boldsymbol{U}_{k+1}^{(1)}] M_{2k}$$

where $(M_{2k})_{1:k+1,1:k} = H_k^{(0)}$ and $(M_{2k})_{k:2k+1,k+1:2k} = H_k^{(1)}$, while the other elements of M_{2k} are equal to zero [22]; see also [3]. Once again, after ℓ cycles, $M_{\ell k}$ will be upper Hessenberg with a block (sparse band) structure.

3.2 | The Nonsymmetric Lanczos Iteration

Given A, c and a so-called shadow vector $\tilde{c}$, the nonsymmetric Lanczos iteration generates two sets of vectors $\{u_i\}_{i \geq 1}, \{w_i\}_{i \geq 1}$ such that $w_i^T u_i = 1$ and $w_j^T u_i = 0$ for $j \neq i$, and are therefore *biorthogonal*. The Lanczos iteration has been used in the literature for eigenvalue computations [27], linear systems [2, Ch. 28], [28, section 3], [29], Model Order Reduction [30–32].

For u_1, w_1 such that $w_1^T u_1 = 1$ (e.g., $u_1 = c/\|c\|$, $w_1 = \tilde{c}/(u_1^T \tilde{c})$), the next iterates are computed by means of the recurrences[2]

$$\tilde{u}_{i+1} = Au_i - u_i \alpha_i + u_{i-1} \beta_i, \qquad \tilde{w}_{i+1} = A^T w_i - w_i \alpha_i - w_{i-1} \delta_i$$

with $\alpha_i = w_i^T A u_i$, $\beta_0 = 0, \delta_0 = 0$, while $\beta_{i+1}, \delta_{i+1}$ scale $\tilde{u}_{i+1}$ and $\tilde{w}_{i+1}$, respectively, so that $w_{i+1}^T u_{i+1} = 1$ [2, section 7.1]. For later convenience, in the following, we consider $\beta_{i+1} = \|\tilde{u}_{i+1}\|$ so that $\|u_{i+1}\| = 1$, and $\delta_{i+1} = \tilde{w}_{i+1}^T u_{i+1}$. In all our experiments, we set $\tilde{c} = c$.

After k iterations, the two recurrences can be compactly written as

$$AU_k = U_k T_k + u_{k+1} t_{k+1,k} e_k^T, \quad A^T W_k = W_k T_k^T + w_{k+1} t_{k+1,k} e_k^T \tag{3.5}$$

where T_k is tridiagonal, and has diagonal elements equal to α_i and off-diagonal elements equal to β_{i+1} (upper part) and δ_{i+1} (lower part). Moreover, $U_k = [u_1, \ldots, u_k]$, $W_k = [w_1, \ldots, w_k]$. Assuming that no breakdown occurs, in exact arithmetic, the matrices U_k, W_k satisfy $W_k^T U_k = I_k$ for all $k > 0$, and no explicit full biorthogonalization needs to be enforced. Hence, unless they are used for later purposes, the two bases do not need to be stored. Biorthogonality also ensures that $T_k = W_k^T A U_k$, which is obtained, for example, by premultiplying the first Lanczos recurrence by W_k^T; a (1)-inverse of U_k is thus available by construction.

Loss of biorthogonality of the generated Lanczos bases is a very well-known plague of the iteration, where breakdown can occur if for some index i, it happens that $w_i^T u_i = 0$. Over the years, several different strategies have been adopted. A common strategy aims at recovering some sort of partial biorthogonality [11, 34], also called semiduality in Reference [35]. Breakdown can, in some cases, be "cured" by a so-called look-ahead procedure, [36, 37], [38]. A less dramatic but still highlighted weakness of the Lanczos recurrence is the fact that the recurrence requires knowing A^T, which may not be available for certain applications. In the following, we shall assume that A^T can be accessed as easily as A, and that no incurable breakdown takes place.

In our context, the Lanczos recurrence becomes particularly appealing because it is inherently three-term, and the matrix T_k is tridiagonal by construction. Hence, as opposed to truncated and restarted procedures, there is no need for a parameter selection k_t and, at least in exact arithmetic, the biorthogonality is preserved. The possibly severe loss of biorthogonality as iterations proceed is due to finite precision computations: Loss of biorthogonality manifests itself both at failing $w_i^T u_i = 1$, but also at failing $w_j^T u_i = 0$, for $j \neq i$. As a result of the latter, in finite precision arithmetic, the method will somehow resemble a truncated method, but without a clear-cut explicit truncation point: Biorthogonality will degrade with the index distance between the basis vectors u_i, w_j. For these reasons, nonsymmetric Lanczos is a good target for sketching strategies to mitigate the aforementioned problems caused by finite precision arithmetic.

4 | Some Spectral Properties of Krylov Decompositions

In this section, we refine certain structural spectral properties of the Rayleigh quotient matrix appearing in the Krylov decomposition. Hence, these properties apply to all methods described in the previous section. As we shall see, they also apply to the subspace embedding procedures that will be described in Section 6, since it will be shown that these procedures also fit the Krylov decomposition framework.

Let us consider the Krylov decomposition (2.1) with the assumptions that B_k is upper Hessenberg and $b_{k+1} \propto e_k$. The approaches recalled in the previous section all fit this setting. It is well known that as soon as the Krylov decomposition deviates from the full Arnoldi relation (i.e., not all columns of U_k are mutually orthonormal), then the spectral properties of the reduced matrix B_k may be distorted [6]. For instance, without orthogonality of the basis, the field of values $\mathcal{W}(B_k)$ may significantly differ from that of A. On the other hand, in the orthonormal case it holds that $\mathcal{W}(B_k) \subset \mathcal{W}(A)$. The set $\mathcal{W}(A)$ is used, for instance, in the analysis of functions of matrices and of their approximations. More generally, for any f analytic in $\mathcal{W}(A)$ it holds that [39]

$$\|f(A)\| \leq (1 + \sqrt{2}) \sup_{z \in \mathcal{W}(A)} |f(z)| \tag{4.1}$$

Hence, the fact that $\mathcal{W}(B_k) \subset \mathcal{W}(A)$ ensures that the bound can also be adopted for $\|f(B_k)\|$.

We next show that for a nonorthogonal basis, certain anomalous spectral behavior of B_k may have very little impact on the performance of the method. This suggests, as already discussed in References [20, 40], that $\mathcal{W}(B_k)$ may not be representative of the recurrence behavior. More precisely, and based on Proposition 4.1 below, it was stated in Reference [20] that a more descriptive region is given by the *effective* field of values, which is defined by means of the projection of the Rayleigh quotient matrix B_k onto the invariant space associated with the eigenvalues not far from the field of values of A. We refer to [20] for details and computational examples.

We first recall a technical result from [20] on the decay of the entries in the unitary matrices of the Schur decomposition of $(\widehat{B}_k + \widehat{g}_k e_k^T)^*$. The result states that if some of the eigenvalues of $(\widehat{B}_k + \widehat{g}_k e_k^T)^*$ are distant from the field of values of $\widehat{B}_k$, then the components of the corresponding Schur vectors will have a decay pattern, so that, in particular, the first component of each Schur vector has a value $\mathcal{O}(\delta^k)$ for some $0 < \delta < 1$. This structural property plays a crucial role in describing the quality of the given a Krylov approximation space.

Proposition 4.1 ([20, Lemma 4.3]). *Let $(\widehat{B}_k + \widehat{g}_k e_k^T)^* = QRQ^*$ be the Schur decomposition of the given matrix, with $\widehat{B}_k$ upper Hessenberg and irreducible. For $\eta > 0$ assume that there exists an index $\bar{\imath} \leq k$ such that $\operatorname{dist}(R_{ii}, \mathcal{W}(\widehat{B}_k)) = O(\eta^{-1})$, $i = 1, \ldots, \bar{\imath}$ for an appropriate spectral ordering in the Schur decomposition. Then it holds that $Q_{1,i} = O((\eta/\xi)^k)$ for $i = 1, \ldots, \bar{\imath}$ with $\xi = \max_{j \leq \bar{\imath}} |1/(\widehat{B}_k)_{j+1,j}|$.*

The following result applies the proposition above to an expanding matrix. In particular, in our setting, the result reads as follows: given the matrix B_{k-1}, if at the kth iteration the matrix is expanded by adding a new row and a new column in a way that the new matrix B_k has an isolated eigenvalue with respect to the rest of the spectrum of both B_k and B_{k-1}, then a decay of the magnitude of the corresponding vector entries in the Schur form can be observed.

Corollary 4.2. *Let $B_k \in \mathbb{R}^{k \times k}$ be irreducible upper Hessenberg, and let B_{k-1} be its principal $(k-1) \times (k-1)$ left upper part. Let $\mathcal{D}_k$ be the smallest between the Gerschgorin disks by row and by column, centered at $(B_k)_{k,k}$. If $\mathcal{D}_k$ does not intersect $\mathcal{W}(B_{k-1})$, then it is possible to write B_k as $B_k = \widehat{B}_k + \widehat{g}_k e_k^T$ and there exists an index $\bar{\imath}$ for which Proposition 4.1 applies.*

Proof. The hypothesis on the Gerschgorin disk ensures that there exists an eigenvalue away from the field of values of the $(k-1)\times(k-1)$ principal part of B_k. We define $\widehat{B}_k$ with $(\widehat{B}_k)_{k,k} = \tau$, so that $B_k = \widehat{B}_k - e_k(\tau - (B_k)_{k,k})e_k^T$. We chose $\tau \in \mathbb{R}$ so as to be well inside $\mathcal{W}(\widehat{B}_k)$. Hence, B_k satisfies the hypotheses of Proposition 4.1 for the eigenvalue contained in the isolated Gerschgorin disk. □

Lemma 4.3 in Reference [20] was used to show that the field of values of B_k may be a misleading information in the context of Sylvester equations: The distorted, outlying eigenvalues of B_k do not contribute to the actual solution of the problem. To appreciate the result of Corollary 4.2 we need to put it into the Krylov subspace context. Assume that B_k is generated by one of the discussed recurrences, so that $\widehat{b}_{k+1} \propto e_k$ in Equation (2.2) and that $\widehat{B}_k$ is upper Hessenberg and irreducible. Let $(\widehat{B}_k + \widehat{g}_k e_k^T)^T = QRQ^*$ be the Schur decomposition of the Rayleigh quotient matrix in Equation (2.2). An element in the Krylov subspace $K_k(A,c)$ can be represented as $\phi_\ell(A)c$ for some polynomial ϕ_ℓ of degree $\ell < k$. Moreover, since $c = \widehat{U}_k e_1\gamma$ (in the following we set $\gamma = 1$) and $\widehat{B}_k + \widehat{g}_k e_k^T$ is upper Hessenberg, explicit computation shows that $\phi_\ell(A)c = \widehat{U}_k\phi_\ell(\widehat{B}_k + \widehat{g}_k e_k^T)e_1$ for $\ell < k$. Assume that Corollary 4.2 is satisfied by some j, and let $Q = [Q_1, q_0]$ be a partitioning of Q so that q_0 has a single column, and $|q_0^* e_1| \le \epsilon$ for some $\epsilon > 0$. The matrix $R = [R_{11}, R_{10}; 0, R_{00}]$ is partitioned accordingly, with R_{00} a scalar. Then

$$\phi_\ell(A)c = \widehat{U}_k[Q_1, q_0]\phi_\ell(R)[Q_1, q_0]^* e_1 \tag{4.2}$$

Hence, since the magnitude of the elements in $\phi_\ell(A)c$ depend on the magnitude of the elements of $\phi_\ell(R)$, the quality of the vector $\phi_\ell(A)c$ as an approximation to some target vector depends on the quality of $\phi_\ell(R)$. In particular, if $\phi_\ell(R_{00})$ is very different (e.g., very large) from what is expected, then the approximation may be affected. On the other hand, in the considered hypotheses, $|q_0^* e_1|$ will be very small, possibly far below the roundoff unit, in exact arithmetic. For the sake of the argument, let us first assume that $q_0^* e_1 = 0$. In this case, it follows that

$$\phi_\ell(A)c = \widehat{U}_k Q_1 \phi_\ell(R_{11}) Q_1^* e_1$$

Hence, the vector $\phi_\ell(A)c$ behaves as if the possibly anomalous value of $\phi_\ell(R_{00})$ did not exist. In other words, if there exists an isolated eigenvalue away from the field of values of B_{k-1} such that $\phi_\ell(R_{00})$ is very large but $q_0^* e_1 = 0$, then that eigenvalue will not contribute to an approximation from the Krylov subspace $K_k(A, v)$. Hence, eigenvalue information alone may be misleading. We remark that the discussion above shows that the approximation space is reduced to the subspace generated by $\widehat{U}_k Q_1$. As a consequence, although the outlying eigenvalue does not contribute, the approximation space shrinks. If in the next few iterations more outlying eigenvalues arise, the effective subspace will not expand, making the approximation stagnate. In conclusion, unless a basis self-correction occurs, the iteration will start to deviate from the orthogonal (ideal) one.

Assume next that $q_0^* e_1$ is nonzero, but that ϵ is of the order of machine precision, so that $q_0^* e_1 = \mathcal{O}(\mathtt{u})$, where $\mathtt{u}$ is the roundoff unit. This is what we expect to occur in finite precision arithmetic[3]. The relation (4.2) indicates that

$$\|\phi_\ell(A)c - \widehat{U}_k Q_1 \phi_\ell(R_{11}) Q_1^* e_1\| \le \|\widehat{U}_k\|\ \|\phi_\ell(R)\begin{bmatrix}0\\ I\end{bmatrix}\|\ \epsilon, \quad \epsilon = \mathcal{O}(\mathtt{u})$$

The actual magnitude of the upper bound depends on how large the matrix norms on the right are. In turn, the second of these norms depends on the selected polynomial ϕ_ℓ, which is problem dependent.

In most cases, we expect $\|\widehat{U}_k\|\ \|\phi_\ell(R)[0; I]\|$ to be significantly smaller than $\mathtt{u}^{-1}$, so that the error behavior will be similar to the case $q_0^* e_1 = 0$. Extreme cases can occur, though, whenever R_{00} is such that $\|\phi_\ell(R)[0; I]\| \ge \mathtt{u}^{-1}$. One such exceptional case is reported in Example 8.1 for the exponential $f(\lambda) = \exp(-\lambda)$ with $\mathcal{W}(A)$ contained in the right-half complex plane, so that $\|f(A)\|$ remains moderate (see (4.1)). In that example, after a certain number k of iterations, the classical Lanczos recurrence has generated a matrix B_k with one very outlying eigenvalue on the left-hand plane, that is far away from the field of values of A where f is well behaved, so that $\|f(B_k)e_1\| \gg \mathtt{u}^{-1}$. As a consequence, the approximation error with respect to the target vector $f(A)c$ increases tremendously. The proposed stabilization procedure aims to cure precisely this type of misbehavior, by transforming the reduced matrix in a way that its spectrum is not too far from $\mathcal{W}(A)$, see Section 7.

In general, this analysis suggests that strategies that are able to correct anomalous spectral behaviors of the nonorthogonal recurrences are desirable.

5 | Randomized Row Selection

Randomization procedures have recently been used in the Krylov subspace context to improve the linear independence of a truncated basis. In this section, we recall some commonly employed randomized row selection strategies and point out certain key subspace properties that will be used in later sections, together with Krylov decompositions.

Given a full column rank matrix $U \in \mathbb{R}^{n\times(k+1)}$ and an integer s with $(k+1) < s \ll n$, we consider the problem of finding a linear combination of s rows of U, denoted as $\mathcal{S}(U)$, with smallest possible condition number in the Euclidean norm.

In our context, where the Krylov subspace basis is generated iteratively, the basis is not all available at the same time, unless it is all stored. With these premises, the selection operator $\mathcal{S}$ has to be chosen a priori, turning to a suboptimal—in some sense—selection.

In the following, we consider two examples of a priori selections, based on randomized subsampling strategies. Accessible descriptions of these and other methods can be found in the References [41, 42]; we follow the derivation in the latter. A common denominator of the approaches we are going to consider is that they are *oblivious* subspace samplings, that is, they are not associated with a specific subspace, but with a class of matrices (basis) satisfying certain properties.

We start by recalling a founding definition for uniform embeddings, in our notation; see, for example, [42, Def. 2.1].

Definition 5.1. A $(1 \pm \varepsilon)\,\ell_2$-subspace embedding for the column space of $U \in \mathbb{R}^{n\times k}$ is a linear operator $\mathcal{S} : \mathbb{R}^n \to \mathbb{R}^s$ such that[4], for all $x \in \mathbb{R}^k$

$$\|\mathcal{S}(Ux)\|_2^2 = (1 \pm \varepsilon)\|Ux\|_2^2$$

We remark that this definition was also employed in [43, def. 1] and recalled in [41, 9.1] for Johnson-Lindenstrauss and uniform embeddings, respectively, whereas for Gaussian embeddings, the relation $\|\mathcal{S}(Ux)\|_2 = (1 \pm \varepsilon)\|Ux\|_2$ is used, for instance, in [41, 8.1].

The linear operator $\mathcal{S}$ can be identified with a matrix $S \in \mathbb{R}^{s\times n}$, so that we can also write $\mathcal{S}(Ux) = SUx$ for any x. As short-hand notation, when we write $\mathcal{S}(U)$ we mean that $\mathcal{S}$ is applied to each column of U. The definition does not depend on the specific basis used for spanning the space. In particular, this implies that U could be taken to have orthonormal columns, so that it holds that $\|\mathcal{S}(Ux)\|_2^2 = (1 \pm \varepsilon)\|x\|_2^2$, from which it also follows

$$\|I_k - (SU)^T SU\|_2 \leq \varepsilon$$

Oblivious embeddings represent a distribution on matrices of dimension $s \times k$, where s depends on n, k, ε and the parameter δ, the failure probability, so that $1 - \delta$ represents the least probability of selecting an $(1 \pm \varepsilon)$ ℓ_2-subspace embedding for a fixed $n \times k$ matrix U drawn from this distribution. Using standard terminology [41, 42], in the following we say that a subspace embedding satisfies a certain property with *high probability* if δ is small.

We also recall a result from Reference [13] for later convenience, where we use the notation

$$\cos\angle(u, v) = \frac{u^T v}{\|u\|\ \|v\|}$$

for the cosine of the angle between two nonzero vectors.

Lemma 5.2 ([13, Lemma 4.3]). *For $\varepsilon > 0$, let u, v satisfy $\|Su\|_2 = (1 \pm \varepsilon)\|u\|_2$ and $\|Sv\|_2 = (1 \pm \varepsilon)\|v\|_2$ for the $(1 \pm \varepsilon)$ ℓ_2-subspace embedding S. Then, with high probability, it holds that*

$$\frac{\cos\angle(u, v) - \varepsilon}{1 + \varepsilon} \leq \cos\angle(Su, Sv) \leq \frac{\cos\angle(u, v) + \varepsilon}{1 - \varepsilon}$$

5.1 | Randomized Subsampling

The idea consists of sampling a small number of rows of U uniformly at random and rescaling by means of a probability measure. This works well for dense matrices, as is our case. We consider S an $s \times n$ portion of the identity matrix, where each nonzero entry per row is given by $1/\pi_i$, and π_i is the probability of that row to be selected. Proving a relation of the type $\|\mathcal{S}(Ux)\|_1 = (1 \pm \varepsilon)\|Ux\|_1$ (in the ℓ_1 norm now) is involved and we refer to [42, section 3.1] and references therein. We will denote this operator as S_{RS} or S_{RS}.

5.2 | Johnson-Lindenstrauss Transform

This may be considered a more costly improvement of the procedure above, as it aims to avoid zero rows in the sampled operator image. In this case, the operator is given by the composition of a diagonal matrix D of ± 1 with probability $1/2$, then a Hadamard or Fourier-type transform H, and then a row sampling operator J, that is, [26, section 9]

$$SU = \sqrt{\frac{n}{s}}\, J H D U$$

The cost of applying this product to a vector u will be $O(n \log n)$ [42, section 2.1]. Then S samples s rows with $s = \Omega(\varepsilon^{-2}(\log k)(\sqrt{k} + \sqrt{\log n})^2)$ and with probability at least 0.99 for a fixed U it holds

$$\|I_k - (SU)^T SU\|_2 \leq \varepsilon$$

We will denote this operator as S_{JL} or S_{JL}.

Other randomization operators can be employed, although not all are efficient in terms of computational costs. For instance, a properly scaled matrix S with normally distributed entries has been considered in the literature, although matrix-vector products with these embedding operators become too expensive within an iterative procedure.

Remark 5.3. If all columns of U were known simultaneously, then the problem could be stated as:

Given $U \in \mathbb{R}^{n\times(k+1)}$, for some s with $k + 1 < s \ll n$, find S giving the best conditioned matrix

$$\widehat{U} := UR^{-1}$$

where $S(U) = QR$ is the QR decomposition of $S(U)$.

This is a combinatorial problem whose computational cost becomes too high for large n. Most importantly, it requires all columns to be available at the same time. In our context, where the Krylov subspace basis is generated iteratively, the solution of the problem above would imply storing the whole Krylov decomposition basis, something we cannot afford.

6 | Subspace Embeddings and Krylov Decompositions

In this section, we report the main results of the paper by closing the gap between subspace embeddings and Krylov decompositions. Recently, it has been shown that reduced sketched bases used within a Krylov iteration satisfy a reduced Krylov decomposition relation; see [40, Pro] and Proposition 6.1 below. Here we show that subspace embedding indeed fully determines a well-defined and well-conditioned Krylov decomposition with high probability, under the hypothesis of linear independence of the original basis.

Assume that a maximum of $\overline{k}$ iterations can be afforded (in the computational experiments, we will denote this value with

$k_{\max}$). Then, consider a tall full column rank matrix $U_{\bar{k}}$ whose columns are a basis for the generated Krylov subspace. Let $SU_{\bar{k}} = Q_{\bar{k}}R_{\bar{k}}$ be the reduced QR decomposition of $SU_{\bar{k}}$. Thanks to the ε-embedding property of the sketching, it holds that (see, e.g., [12])

$$\kappa(U_{\bar{k}}R_{\bar{k}}^{-1}) \leq \sqrt{\frac{1+\varepsilon}{1-\varepsilon}} \tag{6.1}$$

with high probability. In particular, from the definition of subspace embedding, it follows that

$$(1-\varepsilon)\|U_{\bar{k}}R_{\bar{k}}^{-1}v\|^2 \leq \|SU_{\bar{k}}R_{\bar{k}}^{-1}v\|^2 \leq (1+\varepsilon)\|U_{\bar{k}}R_{\bar{k}}^{-1}v\|^2 \tag{6.2}$$

since the embedding property does not depend on the chosen basis. Thus, we will label the quantity $\bar{k}$ as the reference maximum affordable approximation space dimension, for which the probabilistic argument holds.

The bounds (6.1) and (6.2) will be crucial for proving our results, and for assessing the good behavior of the considered methods, with high probability.

The following result provides an Arnoldi-like relation for any row-reduction operator $\mathcal{S}$ applied to (2.1). In particular, this will apply to the subspace embedding operators discussed in Section 5. We stress that deriving an Arnoldi-like relation for the sketched basis turned out to be crucial for the analysis of sketched Krylov subspace methods [20, 40]. The following result is the same as that in [40, Pro] where specific choices of $\mathcal{S}$ and B_k were made. We include its proof because it is constructive.

Proposition 6.1. *Let $\mathcal{S} : \mathbb{R}^n \to \mathbb{R}^s$ be a sketching operator. Let $A \in \mathbb{R}^{n\times n}$ and $[U_k, u_{k+1}] \in \mathbb{R}^{n\times k+1}$ with linearly independent columns. Then for the quantities constituting (2.1) there exist $\widehat{r}$, ρ and $q \perp \mathcal{S}U_k$ such that*

$$\mathcal{S}(AU_k) = \mathcal{S}(U_k)(B_k + \widehat{r}e_k^T) + \rho q b_{k+1,k}e_k^T$$

where $b_{k+1} = b_{k+1,k}e_k$ and e_k has k components.

Proof. Let $SU_{k+1} = Q_{k+1}R_{k+1}$ be the reduced QR decomposition of the given matrix, and $R_{k+1} = [R_k, r; 0, \rho]$, $Q_{k+1} = [Q_k, q]$. Let also $\underline{B}_{k+1} = [B_k; b_{k+1,k}e_k^T]$. We have

$$\begin{aligned} SAU_k &= SU_{k+1}\underline{B}_{k+1} = Q_{k+1}R_{k+1}\underline{B}_{k+1} \\ &= Q_{k+1}\begin{bmatrix} R_kB_k + rh_{k+1,k}e_k^T \\ [0^T, \rho]\underline{B}_{k+1} \end{bmatrix} \\ &= Q_kR_k(B_k + R_k^{-1}rb_{k+1,k}e_k^T) + \rho q b_{k+1,k}e_k^T \\ &= SU_k(B_k + R_k^{-1}rb_{k+1,k}e_k^T) + q\rho b_{k+1,k}e_k^T \end{aligned}$$

and the result follows. □

We notice that in the proof of Proposition 6.1 we used the factorization of the row-reduced matrix $\mathcal{S}(U_{k+1})$, that is $SU_{k+1} = Q_{k+1}R_{k+1}$; in the sketching context this is called "whitening", and it will also be used in the following. Indeed, the Arnoldi-type procedure to be used in practice is a combination of sketching and whitening, where the orthogonalization of the reduced matrix is highly recommended [12, 15, 16]. Numerical experiments in the literature illustrates the importance of this orthogonalization step. In the following, we assume that whitening of the sketched basis is always performed. As in the proof above, we can refer to the relation

$$SAU_kR_k^{-1} = Q_kR_k(B_k + R_k^{-1}rb_{k+1,k}e_k^T)R_k^{-1} + \frac{\rho b_{k+1,k}}{r_{k,k}}qe_k^T \tag{6.3}$$

with $R_{k+1} = [R_k, r; 0, \rho]$ and $r_{k,k} = e_k^TR_ke_k$, and use the matrices

$$\widehat{U}_k := U_kR_k^{-1}, \qquad \widehat{B}_k := R_kB_kR_k^{-1} + rb_{k+1,k}e_k^TR_k^{-1}$$

where, $\widehat{U}_k$ is the *stabilized* basis and $\widehat{B}_k$ is the corresponding projected and restricted representation of A in the generated space. The term "stabilization" is used here in virtue of property (6.1), which holds for $\widehat{U}_k$ with high probability, for the given sketching operator $\mathcal{S}$.

For practical purposes, it is important to stress that the sketching the procedure is performed iteratively as the recurrence progresses, by adding a column to the sketched basis SU_k and expanding the triangular matrix R_k at each iteration, as k grows. This allows us to avoid storing the full basis, while controlling the computational costs.

We next show that the rank-one modification of the reduced matrix B_k leads to an *equivalent* Krylov decomposition of the original (unsketched) Krylov decomposition. This result appears to be new. The formulation of the whitened-sketched procedure in terms of the Krylov decomposition allows us to highlight the advantages of the method in a well-established context, in which all the aforementioned methods can be framed. In particular, this result completes the path of including the sketching strategy into the Krylov framework, first taken in Reference [40] with the Proposition 6.1 above.

Proposition 6.2. *Assume that U_{k+1} is full rank. Let R_{k+1} be the upper triangular matrix of the QR decomposition of SU_{k+1}, with $R_{k+1} = [R_k, r_{k+1}; 0, \rho_{k+1}]$, and $\widehat{U}_{k+1} = [\widehat{U}_k, \widehat{u}_{k+1}] = U_{k+1}R_{k+1}^{-1}$. Then any Krylov decomposition (2.1) can be transformed by sketching and whitening in the following equivalent Krylov decomposition*

$$A\widehat{U}_k = \widehat{U}_k\widehat{B}_k + \widehat{u}_{k+1}\widehat{\beta}_{k+1}e_k^T, \quad \widehat{B}_k = R_kB_kR_k^{-1} + r_{k+1}b_{k+1,k}e_k^TR_k^{-1}, \quad \widehat{\beta}_{k+1} = \rho_{k+1}b_{k+1,k}r_{k,k}^{-1}$$

Proof. Consider the Krylov decomposition $AU_k = U_kB_k + u_{k+1}b_{k+1,k}e_k^T$. Add and subtract $U_kR_k^{-1}rb_{k+1,k}e_k^T$ to have

$$AU_k = U_k(B_k + R_k^{-1}rb_{k+1,k}e_k^T) + (-U_kR_k^{-1}r_{k+1} + u_{k+1})b_{k+1,k}e_k^T$$

Next, with little algebra with R_k^{-1} we obtain

$$\begin{aligned} AU_kR_k^{-1} = U_kR_k^{-1}(R_kB_kR_k^{-1} + rb_{k+1,k}e_k^TR_k^{-1}) \\ + (-U_kR_k^{-1}r_{k+1} + u_{k+1})b_{k+1,k}e_k^TR_k^{-1} \end{aligned}$$

Finally, we observe that $[\widehat{U}_k, \widehat{u}_{k+1}] = [U_k, u_{k+1}]R_{k+1}^{-1}$ so that for the last column it holds that $\widehat{u}_{k+1} = -U_kR_k^{-1}r_{k+1}\rho_{k+1}^{-1} + u_{k+1}\rho_{k+1}^{-1}$. Hence,

$$AU_k R_k^{-1} = U_k R_k^{-1}(R_k B_k R_k^{-1} + r_{k+1} b_{k+1,k} e_k^T R_k^{-1}) + \hat{u}_{k+1}\rho_{k+1} b_{k+1,k} e_k^T R_k^{-1}$$

The result thus follows. □

This sketching-based Krylov decomposition may be interpreted as a way to transform the Krylov recurrence into an equivalent one, with the hope of improving stability. Indeed, for the decomposition to be well posed, and for the recurrence to make progress in expanding the space, it is sufficient that the columns of $[\widehat{U}_k, \hat{u}_{k+1}]$ be linearly independent. We measure this independence in terms of the angle between $\hat{u}_{k+1}$ and $\widehat{\mathcal{U}}_k = \text{Range}(\widehat{U}_k)$. More precisely, for $\|\hat{u}_{k+1}\| = 1$, we define the canonical angle between $\hat{u}_{k+1}$ and $\widehat{\mathcal{U}}_k$ as the quantity

$$\Theta_k(\widehat{U}_k, \hat{u}_{k+1}) = \min_{v \in \widehat{\mathcal{U}}_k, \|v\|=1} \angle(v, \hat{u}_{k+1}) \tag{6.4}$$

where $\angle(v, \hat{u}_{k+1}) := \cos^{-1}|v^T \hat{u}_{k+1}|/(\|v\|\ \|\hat{u}_{k+1}\|)$ [44, p. 45].

The following result shows that linear independence is preserved, in probabilistic terms.

Proposition 6.3. *Let $0 < \varepsilon < 1$. Let $\widehat{U}_{k+1} = [\widehat{U}_k, \hat{u}_{k+1}]$ be a full rank sketched Krylov matrix. Let Θ_k be as defined in Equation (6.4). Then, with high probability,*

$$\cos(\Theta_k) \le \varepsilon$$

Proof. Let V_k be a matrix with orthonormal columns such that $\text{Range}(V_k) = \widehat{\mathcal{U}}_k$. Then $\Theta_k(\widehat{\mathcal{U}}_k, \hat{u}_{k+1}) = \angle(V_k V_k^T \hat{u}_{k+1}, \hat{u}_{k+1})$; [10, section 6.4.3]. The cosine of this angle is given by

$$\cos\angle(V_k V_k^T \hat{u}_{k+1}, \hat{u}_{k+1}) = \frac{(V_k V_k^T \hat{u}_{k+1})^T \hat{u}_{k+1}}{\|\hat{u}_{k+1}\| \|V_k V_k^T \hat{u}_{k+1}\|}$$

Using Lemma 5.2, we have

$$\frac{\cos\angle(V_k V_k^T \hat{u}_{k+1}, \hat{u}_{k+1}) - \varepsilon}{1+\varepsilon} \le \cos\angle(S V_k V_k^T \hat{u}_{k+1}, S\hat{u}_{k+1}) \tag{6.5}$$

We recall that

$$S V_k V_k^T \hat{u}_{k+1} = S\widehat{U}_k T_k^{-1} V_k^T \hat{u}_{k+1} = Q_k T_k^{-1} V_k^T \hat{u}_{k+1}$$

and $S\hat{u}_{k+1} = q_{k+1}$. Then $\cos\angle(S V_k V_k^T \hat{u}_{k+1}, S\hat{u}_{k+1}) = \cos\angle(Q_k T_k^{-1} V_k^T \hat{u}_{k+1}, q_{k+1}) = 0$. Plugging this into (6.5) we finally obtain

$$\cos\angle(V_k V_k^T \hat{u}_{k+1}, \hat{u}_{k+1}) \le \varepsilon \tag{6.6}$$

□

As a consequence of the angle result above, we observe that as long as linear independence is maintained via sketching, the space keeps growing and the approximation keeps making progress. This is in line with previous discussions for truncated methods [9]. In later Figure 4(right) (cf. Section 8) we report on the convergence behavior of the nonsymmetric Lanczos method as s varies for a particular example. The cosine of the angle associated with the space expansion remains smaller than one for all selections of s, with smaller values for s larger, as intuitively expected. In particular, we notice that convergence (Figure 4(left)) is not at all affected by the choice of s, at least for this example.

Finally, we emphasize that if a "naturally truncated" recurrence is employed, such as the Lanczos method, the sketching operation does not have to act as a convergence-enforcing procedure, but mainly as a stabilization strategy, to make the basis better conditioned. Truncation parameters are no longer involved in the convergence behavior. In this context, the size of s can be dictated merely by memory constraints for storing the matrix SU_k, as long as SU_k remains a tall matrix. In other words, s could be chosen as large as the memory allocations for SU_k allow.

The following remark provides a different and new—to the best of our knowledge—interpretation of the role of the sketching operators when the original basis is numerically rank-deficient.

Remark 6.4. In exact arithmetic, the full rank matrices U_k and $\widehat{U}_k$ span the same space, which corresponds to $\mathcal{K}_k(A, c)$. If U_k is not numerically full rank, thanks to (6.1) the matrix $\widehat{U}_k$ will be better conditioned, with high probability. On the other hand, numerically it will partially build a different subspace than a Krylov subspace. Indeed, let $U_k = \widetilde{U}_k M_k$ be the reduced QR decomposition of the constructed nonorthogonal basis U_k. In general, $\text{Range}(U_k) \subseteq \text{Range}(\widetilde{U}_k)$, and equality holds numerically only if M_k is numerically nonsingular. Hence, if U_k is numerically rank deficient, then numerically $\text{Range}(U_k) \subset \text{Range}(\widetilde{U}_k)$ strictly. Now, sketching ensures with high probability that the matrix $\widehat{U}_k$ is well conditioned. Since

$$\widehat{U}_k = U_k R_k^{-1} = \widetilde{U}_k (M_k R_k^{-1})$$

this implies that $\text{Range}(\widehat{U}_k) = \text{Range}(\widetilde{U}_k)$, where $\text{Range}(\widetilde{U}_k)$ may contain spurious vectors, i.e, vectors that do not belong to $\mathcal{K}_k(A, c)$. How faithful $\text{Range}(\widetilde{U}_k)$ is to the original Krylov space for the given data depends on how much information is retained numerically. In Figure 6 of Example 9.2 we provide an illustration of this phenomenon: we show that in the truncated Arnoldi recurrence, if the value of the truncation parameter k_t is not large enough, important information of $\mathcal{K}_k(A, c)$ is lost early during the recurrence, to the point that the sketched iteration cannot recover.

7 | Spectral Distance From the Field of Values of A

In this section, we deepen our understanding of the spectral properties of the Rayleigh quotient matrix B_k by making a closer connection with the field of values of A, $\mathcal{W}(A)$. Thanks to the discussion in Section 4 and the bound in Equation (4.1), we expect this connection to give indications on when a sketching procedure will be successful in stabilizing the underlying algorithm.

In the analysis of orthogonalization-based Krylov subspace methods, the fact that the field of values of B_k is contained in that of A is indeed crucial. In our more general setting, the nonorthogonality of the matrix $[U_k, u_{k+1}]$ in the Krylov decomposition (2.1) may lead B_k to have spectral properties quite different from those of A. It is very easy to construct examples of decompositions where

some of the eigenvalues of B_k fall far beyond the field of values of A. In Example 8.1 below, we report one such instance, for the decomposition associated with the nonsymmetric Lanczos method.

The following result aims to bound the distance of the spectrum of B_k from the field of values of A in terms of the quantities available in the Krylov decomposition, as a means to quantify the distortion caused by the nonorthogonality of the basis. Then the bound will be specialized to the Krylov decomposition stemming from the whitened-sketched procedure. We mention that the geometric properties of Krylov subspaces have long been employed to characterize the quality of Krylov approximations for various problems, see, for example, [45]. The first result holds for any pair of matrices $[U_k, u_{k+1}]$, B_k and nonzero vector b_{k+1} satisfying the conditions of a Krylov decomposition.

Proposition 7.1. *Consider the Krylov decomposition (2.1). Let λ be an eigenvalue of B_k, and let $U = \widetilde{U}_k M_k$ be the reduced QR decomposition of U_k. Then there exists a unit norm vector $y \in \mathbb{C}^n$ such that*

$$|\lambda - y^* A y| \le |\cos(\Theta_k)|\ \|u_{k+1}\|\ \|b_{k+1}^T M_k^{-1}\| \le |\cos(\Theta_k)| \|M_k^{-1}\|\ \|b_{k+1}\|$$

where Θ_k is the canonical angle between Range(U_k) and Range(u_{k+1}).

Proof. In the proof, we drop subscripts. We first recall that in the Krylov decomposition, the columns of $[U, u]$ are linearly independent. Let $U = \widetilde{U} M$ be the reduced QR decomposition of U, with $\widetilde{U}$ having orthonormal columns. By similarity transformation we thus write the Krylov decomposition (2.1) as

$$A\widetilde{U} = \widetilde{U}\widetilde{B} + ub^T M^{-1} \tag{7.1}$$

Then we define $\widetilde{u}$ to make u orthogonal to $\widetilde{U}$, that is $\widetilde{u} = u - \widetilde{U}g$ with $g = \widetilde{U}^T u$. By translation we then transform (7.1) into

$$A\widetilde{U} = \widetilde{U}(\widetilde{B} + gb^T M^{-1}) + \widetilde{u}b^T M^{-1} \tag{7.2}$$

From Equation (7.2) we obtain

$$\widetilde{U}^T A\widetilde{U} = \widetilde{B} + gb^T M^{-1}$$

Let $(\lambda, \widetilde{x})$ be an eigenpair of $\widetilde{B}$ with unit norm $\widetilde{x}$. Note that due to the similarity transformation, λ is also an eigenvalue of B. Then we have

$$\widetilde{x}^* \widetilde{U}^T A \widetilde{U} \widetilde{x} = \lambda + \widetilde{x}^* g b^T M^{-1} \widetilde{x}$$

so that, setting $y = \widetilde{U}\widetilde{x}$ we have $|\lambda - y^* A y| \le \|g\|\ \|b^T M^{-1}\|$. Noticing that $\|g\| = |\cos(\Theta)|\ \|u\|$ the result follows. □

The bound is nontrivial only if $\lambda \notin \mathcal{W}(A)$, and it is generally not sharp, as there may be points in $\mathcal{W}(A)$ that are closer to λ than that obtained in the proof of Proposition 7.1. However, for $\lambda \notin \mathcal{W}(A)$ it provides important information on the quantities that may influence the location of the eigenvalues of B_k with respect to $\mathcal{W}(A)$. More precisely, Θ_k describes the linear independence between U_k and u_{k+1}, while $\|M_k^{-1}\|$ accounts for the conditioning of the basis U_k.

It is crucial to realize that Proposition 7.1 holds for any equivalent Krylov decomposition, whenever U_k and B_k are replaced by the appropriate quantities. This allows us to make similar deductions for all the methods we have considered. Clearly, for a recurrence leading to a Krylov decomposition with $\cos(\Theta_k)$ close to one and large $\|M_k^{-1}\|$, some of the eigenvalues of the reduced matrix B_k may happen to lie away from $\mathcal{W}(A)$. More importantly, if the recurrence is such that both quantities have modest values, the corresponding iteration will behave as if the eigenvalues of B_k belonged to $\mathcal{W}(A)$. This is indeed the case for the sketched iteration, as shown in the next result.

Proposition 7.2. *Consider the sketched-and-whitened Krylov decomposition $A\widehat{U}_k = \widehat{U}_k \widehat{B}_k + \widehat{u}_{k+1}\widehat{\beta}_{k+1} e_k^T$ as in Proposition 6.2. Let λ be an eigenvalue of $\widehat{B}_k$. Then, with high probability, there exists a unit norm vector $y \in \mathbb{C}^n$ such that*

$$|\lambda - y^* A y| \le \frac{\sqrt{1+\varepsilon}}{\sqrt{1-\varepsilon}} \varepsilon\ \|A\|$$

Proof. The bound of Proposition 6.2 for the sketched quantities reads

$$|\lambda - y^* A y| \le |\cos(\Theta_k)|\ \|\widehat{u}_{k+1}\|\ \|\widehat{\beta}_{k+1} e_k^T M_k^{-1}\|$$

where Θ_k is the canonical angle between Range$(\widehat{U}_k)$ and $\widehat{u}_{k+1}$. From Proposition 6.3 we have that $|\cos(\Theta_k)| \le \varepsilon$ and from Definition 5.1, $\|\widehat{u}_{k+1}\| \le 1/\sqrt{1-\varepsilon}$. To bound the other quantities, we need to recover the sketched relation. Proposition 6.2 gives

$$A\widehat{U}_k = \widehat{U}_k \widehat{B}_k + \widehat{u}_{k+1}\widehat{\beta}_{k+1} e_k^T$$

Using the same notation as in the proof of Proposition 7.1, let $\widehat{U}_{k+1} = \widetilde{U}_{k+1} M_{k+1}$ be the reduced QR decomposition of $\widehat{U}_{k+1}$, with $M_{k+1} = [M_k, *; 0, M_{k+1,k+1}]$. Then for $g = \widetilde{U}^T \widehat{u}$,

$$A\widetilde{U}_k = \widetilde{U}_k(\widetilde{B}_k + g\widehat{\beta}_{k+1} e_k^T M_k^{-1}) + \widetilde{u}_{k+1} M_{k+1,k+1} \widehat{\beta}_{k+1} e_k^T M_k^{-1}$$

Multiplying from the left by $\widetilde{u}_{k+1}^T$ we obtain $M_{k+1,k+1}\widehat{\beta}_{k+1} e_k^T M_k^{-1} = \widetilde{u}_{k+1}^T A \widetilde{U}_k$. Hence, $\|\widehat{\beta}_{k+1} e_k^T (\widehat{U}_k^T \widehat{U}_k)^{-1} \widehat{U}_k^T\| = \|\widehat{\beta}_{k+1} e_k^T M_k^{-1}\| = \|M_{k+1,k+1}^{-1} \widetilde{u}_{k+1}^T A \widetilde{U}_k\|$. Thus

$$\|\widehat{\beta}_{k+1} e_k^T M_{k+1}^{-1}\| \le |M_{k+1,k+1}^{-1}|\ \|A\| \le (1+\varepsilon)^{1/2} \|A\|$$

The last inequality follows from the fact that $|M_{k+1,k+1}^{-1}| \le \|M_{k+1}^{-1}\| = \frac{1}{\sigma_{\min}(\widehat{U}_{k+1})}$, and $(1+\varepsilon)^{-1/2} = (1+\varepsilon)^{-1/2}\sigma_{\min}(S\widehat{U}_{k+1}) \le \sigma_{\min}(\widehat{U}_{k+1})$, so that $\|M_{k+1}^{-1}\| \le (1+\varepsilon)^{1/2}$. The inequality thus follows. □

8 | Computational Illustration: Approximating the Action of a Matrix Function

Numerical experiments showing the good behavior of the sketched procedure have been abundantly reported in the recent literature of matrix function evaluations—including the inverse, of eigenvalue problems, and matrix equations, see the references

in the introduction. In this section, we illustrate our theoretical results using the approximation of the action of a matrix function to a vector, while in Section 9 we focus on the solution of the Lyapunov equation, as a model problem. In both cases, we mainly focus on the use of the nonsymmetric Lanczos method as an underlying iteration, which provides a different perspective, while highlighting the regularization properties of the randomization-based procedure.

In the following we mainly compare $f(A)c$ with the vanilla Lanczos approximation $f_k := U_k f(B_k)e_1\beta_0$, where $\beta_0 = \|u_1\|$, and with its stabilized-via-sketching version $\widehat{f}_k := \widehat{U}_k f(\widehat{B}_k)e_1\beta_0$, where $\widehat{U}_k, \widehat{B}_k$ are defined as in Proposition 6.2, and $\beta_0 = (\mathcal{S}(u_1))^T\mathcal{S}(c)$. Without loss of generality, we assume $\|c\| = 1$. In some of the examples, we will also use the truncated Arnoldi basis and its sketched variant, as a well-exercised equivalent Krylov decomposition.

We stress that in a practical implementation, the whole of $\widehat{U}_k$ will not be available, but a two-pass strategy should be employed to recover the final approximation at termination; see, for example, [40, 46] and references therein.

Example 8.1. We work with $f(\lambda) = \exp(-\lambda)$, and consider the matrix $A \in \mathbb{R}^{n\times n}$ with $n = 4900$ stemming from the discretization by finite differences of the operator $\mathcal{L} = \mathcal{L}(u) = -(\exp(-xy)u_x)_x - (\exp(xy)u_y)_y - 100(x+y)u_y + 500u$ in the unit square; c has all equal entries, normalized to have unit norm. The matrix A has eigenvalues all in the right half complex plane. This coarse discretization is used to be able to explicitly compute the vector $f(A)c$ for comparison purposes. The convergence of the Lanczos approximation, in terms of the error norms $\|f(A)c - f_k\|$, $\|f(A)c - \widehat{f}_k\|$ is reported in Figure 1 (dash-dotted and solid lines, resp.) as k grows. The approximation $\widehat{f}_k$ was obtained by means of the stabilized strategy using $\mathcal{S}_{RS}$, the simple randomized subsampling with π_i a normalized random probability from a uniform distribution[5], and $s = 2k_{\max} = 80$.

We observe the occurrence of the error peak at the 10th iteration in the standard iteration, corresponding to a sizable negative eigenvalue (Ritz value) of B_k (of the order of -200), while all other eigenvalues are on the right-hand side of the complex plane. Our analysis in Section 4 implies that the contribution of the undesired eigenvalue to the approximation should be irrelevant. More precisely, let $B_k = QRQ^*$ denote the Schur decomposition of B_k, with $Q = [Q_1, q_0]$ with q_0 consisting of a single vector, and $|q_0^* e_1| \le \epsilon$. According to the discussion around (4.2), we notice that at iteration 10, although $q_0^* e_1 = O(10^{-15})$, the term $f(R_{00})$ satisfies $f(R_{00}) \approx \exp(-(-200)) = O(10^{85})$ so that the quantity $f(R_{00})q_0^* e_1$ is no longer negligible. This gives the peak in the error norm, thus strongly affecting the computation.

In the Lanczos process B_k is tridiagonal, and at $k = 10$ the diagonal value $(B_k)_{k,k} = w_k^T A u_k$ is about -200, while the entries $(B_k)_{k,k-1}$ and $(B_k)_{k-1,k}$ are moderate, leading to an isolated Gerschgorin disk. The large value of $w_k^T A u_k$ is due to a small value of $\widetilde{w}_k^T \widetilde{u}_k$, which scales the new basis vectors $\widetilde{w}_k, \widetilde{u}_k$ so has to have $w_k^T u_k = 1$. Apparently, u_k is close to U_{k-1} and $\widetilde{w}_k$ is almost orthogonal to u_k. Forcing the unit norm of these two vectors results in the shown peak. This irregular convergence behavior is well known and often reported in the literature, and may have contributed to disregarding the Lanczos recurrence for matrix function evaluation purposes.

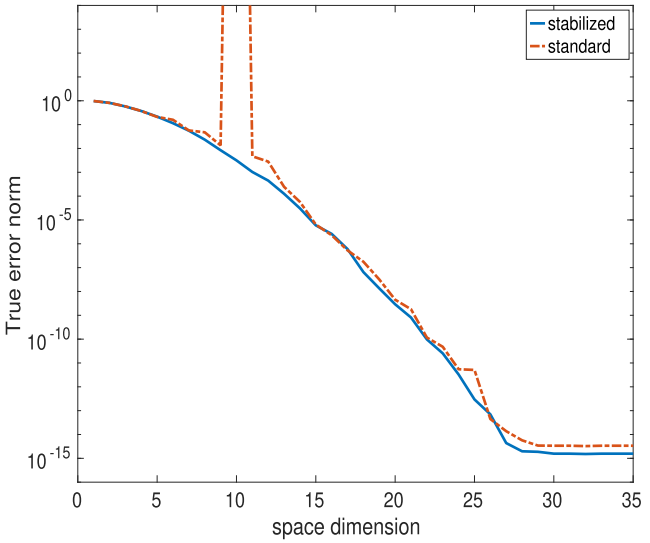

FIGURE 1 | Example 8.1. Error norm—as space dimension increases—between $f(A)c$ and its approximation: Lanczos vector $f_k = U_k f(B_k)e_1$ (dashed) and stabilized (sketched) vector $\widehat{f}_k = \widehat{U}_k f(\widehat{B}_k)e_1\beta_0$ (solid).

The considered simple stabilization strategy relaxes the biorthogonality by making the basis $[\widehat{U}_k, \widehat{u}_{k+1}]$ more linearly independent. This allows the resulting Ritz values to be closer to the good half complex plane, as predicted by Proposition 7.2. For completeness, we mention that the full Arnoldi procedure on this problem behaves exactly like the stabilized procedure.

We explicitly remark that the previous example may serve as a source of inspiration also in the eigenvalue problem setting, where the construction of the left basis W_k can be used to approximate the left eigenvector. In our setting, such construction may be regarded as a parameter-free way to impose a weak orthogonalization property.

Example 8.2. We consider again the same matrix as in the previous example and the same dimension. This time, we analyze the performance when dealing with the function $f(x) = \exp(-\sqrt{A})c$ and c is a normalized random vector. In the left plot of Figure 2 we report the error norm history for the Lanczos iteration (dashed line) and its sketched regularization (solid curve) using $\mathcal{S}_{JL}$ and $s = 2k_{\max} = 360$. In the right plot, we show the angle of the cosine between U_k and u_{k+1} (dashed line) and between $\widehat{U}_k$ and $\widehat{u}_{k+1}$ (solid line) (the left Lanczos sequence pair W_k and w_{k+1} behaves similarly to the right one, and is thus not reported). As expected, the stabilized approach is better able to preserve linear independence (smaller angle cosine) than the standard procedure.

For completeness, corresponding plots for truncated-and-sketched Arnoldi are reported in Figure 3, where the truncation value $k_t = 3$ was used. The full correspondence between the loss

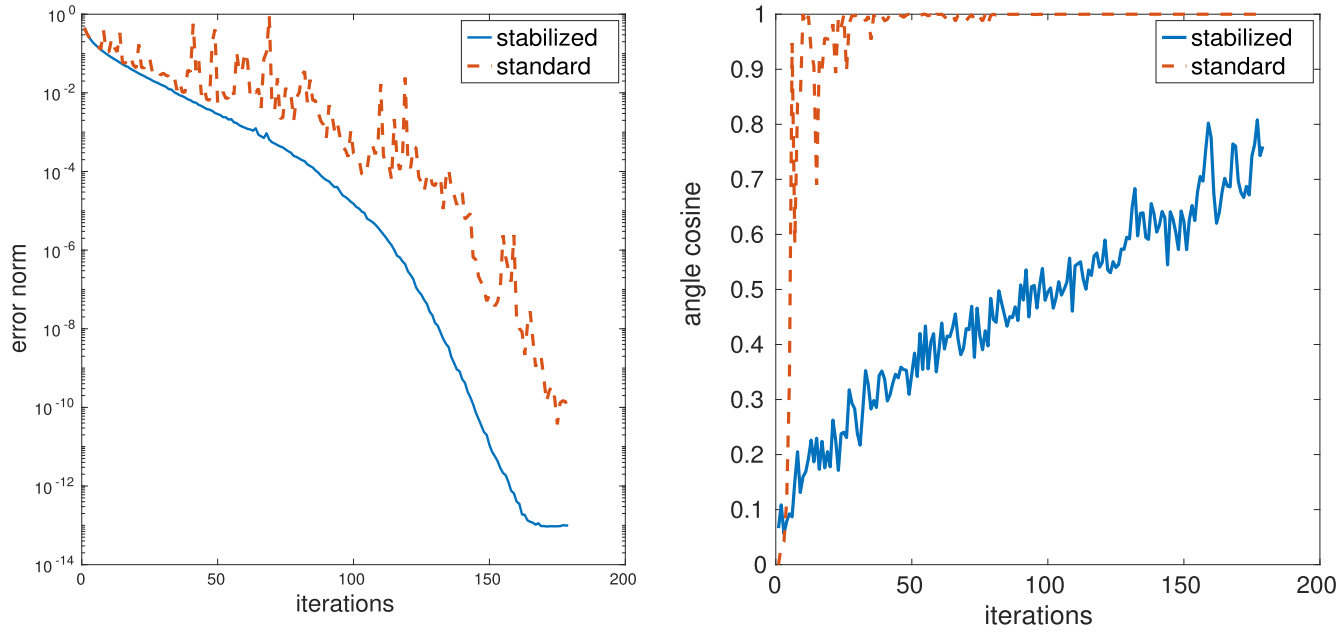

FIGURE 2 | Example 8.2. Left: Error norm—as space dimension increases—between $f(A)c$ and its approximation: Lanczos vector $U_k f(B_k)e_1$ (dashed line) and stabilized (sketched) vector $\widehat{U}_k f(\widehat{B}_k)e_1\beta_0$ (solid line). Right: Cosine of the angle between the Lanczos quantities Range(U_k) and space(u_{k+1}) (dashed line) and between their sketched counterparts Range$(\widehat{U}_k)$ and space$(\widehat{u}_{k+1})$ (solid line).

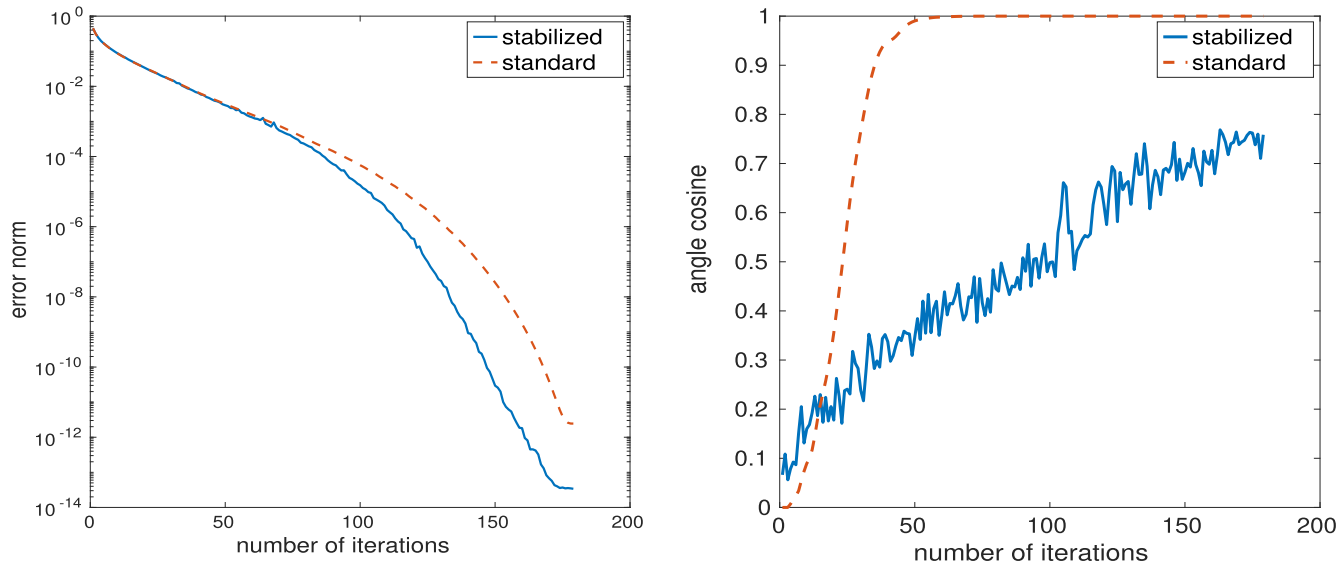

FIGURE 3 | Example 8.2. Left: Error norm—as space dimension increases—between $f(A)c$ and its approximation: Truncated ($k_t = 3$) Arnoldi vector $U_k f(B_k)e_1$ (dashed line) and stabilized (sketched) vector $\widehat{U}_k f(\widehat{B}_k)e_1\beta_0$ (solid line). Right: Cosine of the angle between the Arnoldi quantities Range(U_k) and space(u_{k+1}) (dashed line) and between their sketched counterparts Range$(\widehat{U}_k)$ and space$(\widehat{u}_{k+1})$ (solid line).

of independence in the basis and the degradation of the error norm can be appreciated.

Finally, in Figure 4(left) we report the behavior of nonsymmetric Lanczos as s varies from $s = 190$ to $s = 390$, which appears to be rather insensitive to s, at least on this example. Moreover, as expected, we see in Figure 4(right) that the cosine of the angle between $\widehat{U}_k$ and $\widehat{u}_{k+1}$ is larger for smaller s (upper curves), as less independence is induced by the sketched orthogonality. Nonetheless, independence is preserved in all cases, envisioning a similar convergence history.

9 | Computational Illustration: The Lyapunov Equation in the Lanczos Bases

We consider using the Lanczos bases with sketched regularization for approximating the solution to the following Lyapunov equation

$$AX + XA^T = cc^T$$

Using the two Lanczos recurrences, we can look for an approximation $X_k = U_k Y_k U_k^T \approx X$, and impose the Petrov-Galerkin condition on the residual, using the basis for the left space,

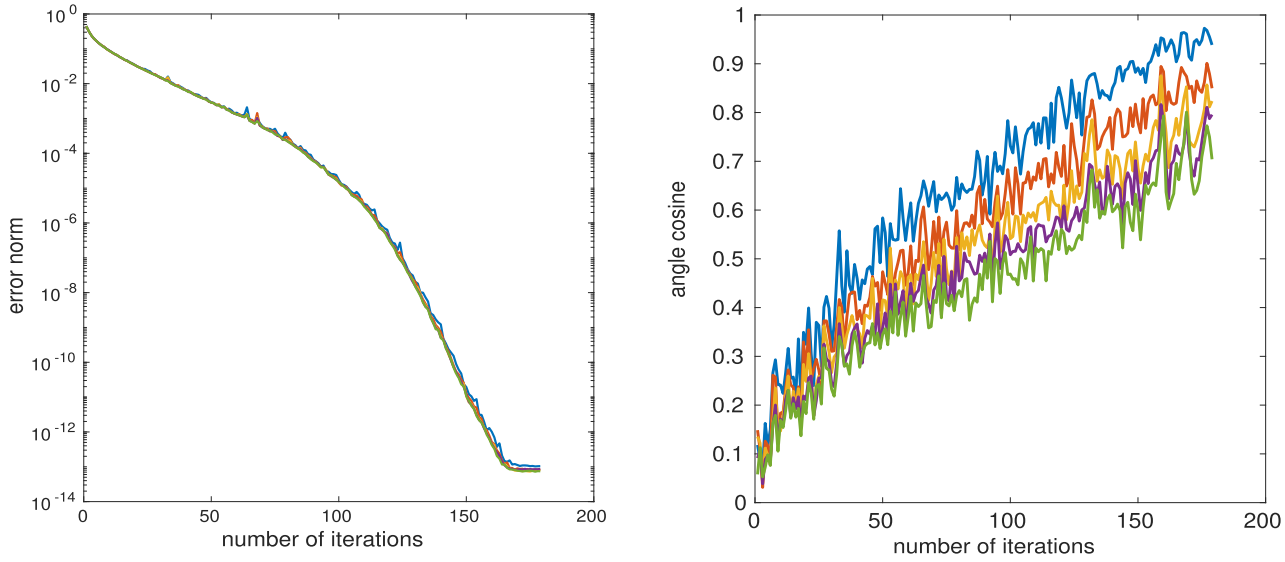

FIGURE 4 | Example 8.2. Exploring different values of the sketching parameter s, $s = 190 : 50 : 390$. Left: Error norm—as space dimension increases—between $f(A)c$ and stabilized (sketched) nonsymmetric Lanczos approximation $\widehat{U}_k f(\widehat{B}_k)e_1\beta_0$. Right: Cosine of the angle between Range($\widehat{U}_k$) and space($\widehat{u}_{k+1}$) for different values of s (the larger s the lower the curve).

W_k. More precisely, let $Res_k = AX_k + X_kA^T - cc^T$. Imposing the condition

$$W_k^T Res_k W_k = 0$$

and using the biorthogonality property $W_k^T U_k = I_k$, we obtain

$$0 = W_k^T AU_k Y_k U_k^T W_k + W_k^T U_k Y_k U_k^T A^T W_k - W_k^T U_k e_1 e_1^T U_k^T W_k \quad (9.1)$$

$$= T_k Y_k + Y_k T_k^T - e_1 e_1^T \quad (9.2)$$

This is what we would get by using the "canonical" Lanczos bases, as it is common in dynamical system reduction processes [30, Part]. We notice in passing that for stability purposes, the biorthogonality property is often relaxed to $W_k^T U_k = D_k$, with D_k diagonal. In the matrix equation context, it is important to realize that the spectral properties of the reduced matrix T_k affect even the existence and uniqueness of the reduced solution matrix Y_k, as the eigenvalues λ_j of T_k should satisfy $\lambda_i \neq -\lambda_j$ for all $i \neq j$ [47]. If $\mathcal{W}(A)$ is, for example, in $\mathbb{C}^+$ then spec$(T_k) \subset \mathbb{C}^+$ is a sufficient condition for preserving a uniquely well defined Y_k. The sketched recurrence helps achieve this goal, as Proposition 7.2 suggests. Any equivalent Krylov decomposition for which a sketched relation, as in Proposition 6.2, holds can be used to obtain an approximation. In the following, we recall the general strategy, also discussed in Reference [20], keeping in mind that it will be applied to the Lanczos recurrence.

Using the regularized basis $\widehat{U}_k = U_k R_k^{-1}$ with the notation in Proposition 6.2, and setting $\widehat{X}_k = \widehat{U}_k \widehat{Y}_k \widehat{U}_k^T$, we can write the residual as

$$\begin{aligned} Res_k &= A\widehat{U}_k\widehat{Y}_k\widehat{U}_k^T + \widehat{U}_k\widehat{Y}_k\widehat{U}_k^T A^T - \widehat{U}_k e_1\widehat{\beta}_0^2 e_1^T\widehat{U}_k^T \\ &= (\widehat{U}_k\widehat{B}_k + \widehat{u}_{k+1}\widehat{\beta}_{k+1}e_k^T)\widehat{Y}_k\widehat{U}_k^T + \widehat{U}_k\widehat{Y}_k(\widehat{U}_k\widehat{B}_k + \widehat{u}_{k+1}\widehat{\beta}_{k+1}e_k^T)^T - \widehat{U}_k e_1\widehat{\beta}_0^2 e_1^T\widehat{U}_k^T \\ &= [\widehat{U}_k, \widehat{u}_{k+1}] \begin{bmatrix} \widehat{B}_k\widehat{Y}_k + \widehat{Y}_k\widehat{B}_k^T - e_1\widehat{\beta}_0^2 e_1^T & \widehat{\beta}_{k+1}\widehat{Y}e_k \\ \widehat{\beta}_{k+1}e_k^T\widehat{Y} & 0 \end{bmatrix} [\widehat{U}_k, \widehat{u}_{k+1}]^T \end{aligned}$$

where $\widehat{\beta}_0 = \|S(c)\|$.

The reduced solution matrix $\widehat{Y}_k$ can be obtained by solving the inner matrix equation $\widehat{B}_k\widehat{Y}_k + \widehat{Y}_k\widehat{B}_k^T = e_1\widehat{\beta}_0^2 e_1^T$. Once again, we recall that $\widehat{U}_k$ is not available, and that the final approximation $\widehat{X}_k$ can be obtained, in factored form, via a two-pass strategy. Further discussion for the truncated and sketched Arnoldi methods can be found in Reference [20].

Remark 9.1. In the derivation above the basis W_k is only used to impose the biorthogonality of U_k, but carries no biorthogonality relations over $\widehat{U}_k$. Sketching could in fact be applied to both bases, that is

$$\widehat{U}_k = U_k R_k^{-1}, \qquad \widehat{W}_k = W_k G_k^{-1}$$

where R_k and G_k are the matrices enforcing biorthogonality of the sketched matrices $S(U_k)$ and $S(W_k)$, see, for example, [38]. The approximate solution is thus written as $X_k = \widehat{U}_k\widehat{Y}_k\widehat{U}_k^T$, while the Petrov-Galerkin condition can again be imposed as

$$\widehat{W}_k^T Res_k \widehat{W}_k = 0$$

Proceeding as in Equation (9.1), we obtain

$$0 = \widehat{W}_k^T A\widehat{U}_k\widehat{Y}_k\widehat{U}_k^T\widehat{W}_k + \widehat{W}_k^T\widehat{U}_k\widehat{Y}_k\widehat{U}_k^T A^T\widehat{W}_k - \widehat{W}_k^T\widehat{U}_k e_1 e_1^T\widehat{U}_k^T\widehat{W}_k \quad (9.3)$$

$$= \widehat{T}_k\widehat{Y}_k K_k^T + K_k\widehat{Y}_k\widehat{T}_k^T - K_k e_1 e_1^T K_k^T \quad (9.4)$$

where $\widehat{T}_k = \widehat{W}_k^T A\widehat{U}_k$ and $K_k = \widehat{W}_k^T\widehat{U}_k$. The biorthogonality of the sketched basis will possibly make K_k well conditioned, though the basis $\widehat{U}_k$ may lose the good properties outlined in the previous sections. Our numerical experiments did not give us enough positive insight to encourage us into pursuing this approach further.

The following numerical example illustrates the behavior of the sketched Lanczos approach, compared to the sketched truncated Arnoldi procedure, for a particularly difficult problem. The sketching operator S_{JL} was used throughout.

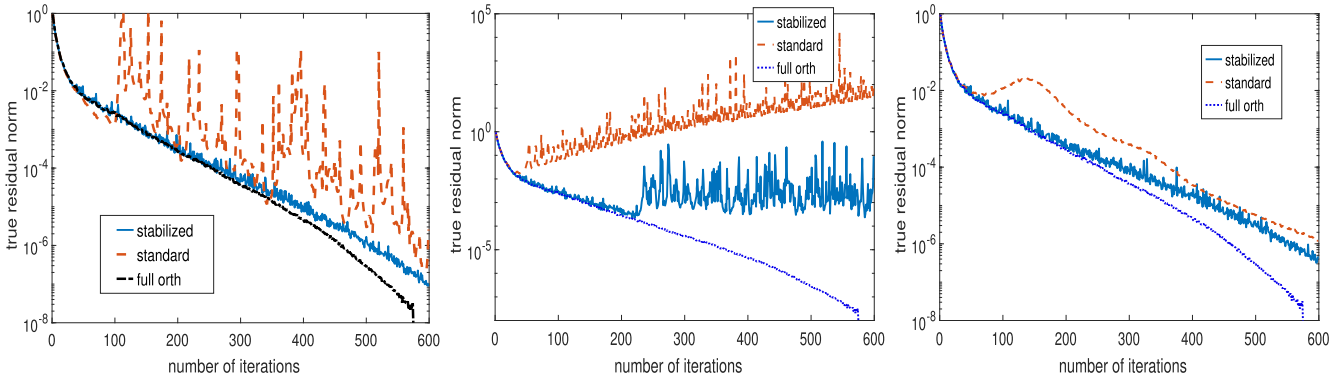

FIGURE 5 | Example 9.2. Residual Frobenius norm for standard methods (dashed line) and stabilized via sketching methods (solid line). The full orthogonalization recurrence is included as a reference method in all plots. Left: Nonsymmetric Lanczos. Middle: Truncated Arnoldi, $k_t = 5$. Right: Truncated Arnoldi, $k_t = 10$.

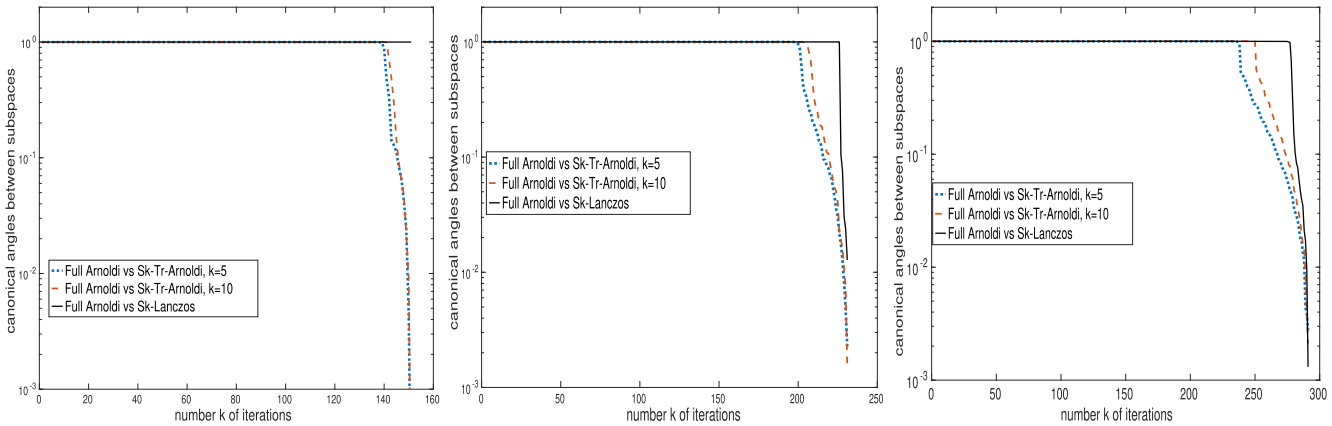

FIGURE 6 | Example 9.2. Cosine of all canonical angles between the full Arnoldi basis $\boldsymbol{U}_k$ and the corresponding sketched bases after k iterations. Left: $k = 150$. Middle: $k = 230$. Right: $k = 290$.

Example 9.2. We consider the matrix $A \in \mathbb{R}^{n\times n}$ with $n = 16\,641$ stemming from the discretization by finite elements of the operator $\mathcal{L}(u) = -\epsilon\Delta u + 2y(1-x^2)u_x + 2x(1-y^2)u_y$ (convection diffusion with recirculating wind [48]) with $\epsilon = 0.1$, on the unit square with homogeneous boundary conditions. This is a hard problem from the benchmark set in the IFISS[6] software [49], and a regularized SUPG discretization on a stretched grid was used to generate the matrix. The vector c was chosen to have normally distributed random entries. The left plot of Figure 5 reports the convergence of the original and sketched Lanczos methods applied to the corresponding Lyapunov equation. In all cases, $s = 2k_{\max} = 1200$ was used. The regularizing effect of the reduction procedure in the Lanczos recurrence can be fully appreciated in the solid curve. The remaining plots of Figure 5 show the convergence history of the original and sketched Arnoldi methods, for the truncation parameter $k_t = 5$ (middle) and $k_t = 10$ (right). The convergence obtained by the full orthogonalization recurrence is included in all plots as the reference method. Clearly, it is very difficult to know a priori which truncation value will ensure good convergence of the Arnoldi procedure. On the other hand, the Lanczos approach completely avoids this issue. For this particular case, it seems that if k_t is too small, too much information gets lost in the generated space, and the sketched Arnoldi approach is unable to recover (middle plot). According to Remark 6.4, we deepen the analysis of this phenomenon by comparing the space built by the sketched bases $\widehat{U}_k$ and the full Arnoldi basis $\boldsymbol{U}_k$; here $\widehat{U}_k$ is generated by either the sketched truncated Arnoldi method with $k_t = 5, 10$ or by the sketched Lanczos recurrence. The comparison is done by computing the cosine of all canonical angles between the two spaces[7]. As long as the original matrix U_k is full rank, we expect all cosines to be equal to one, or very close to it. If U_k is numerically rank deficient, then $\widehat{U}_k$ will be better conditioned, with high probability, but it will partially build a different subspace than a Krylov subspace. Figure 6 reports the cosines of the canonical angles at iteration $k = 150, 230, 290$, showing that for all k considered, the Sketched Lanczos iteration is capable of retaining far more information of $\boldsymbol{U}_k$ than the other bases.

The previous example highlights the fact that the weak implicit biorthogonalization of the Lanczos method overcomes the possible misconvergence of truncation based procedures, in case a too small truncation parameter is employed.

10 | Conclusions and Outlook

Sketching strategies have been used in truncated Arnoldi recurrence-based methods for a variety of numerical linear algebra problems, although only very recently their algebraic

properties have started to be investigated. We have shown that a larger class of methods, verifying a *Krylov decomposition* relation can be analyzed using the same framework, and we have derived an equivalent Krylov decomposition that includes sketched quantities. The class of methods comprises recurrences that have inherently few terms, so that no truncation needs to be imposed. In this context, we have proved that randomization can help stabilize the recurrence by ensuring more linearly independent iterates, and by controlling the deviation of the eigenvalues of the reduced matrix from the target spectral region of A. On the other hand, we have highlighted that sketching may create spurious information if the original basis loses numerical rank too soon. We have mainly focused on the performance of the nonsymmetric Lanczos recurrence, which satisfies a biorthogonalization property in exact arithmetic. Although in finite precision arithmetic, this property no longer holds, our (albeit limited) experimental experience has shown that the "quasi" biorthogonality, together with the sketched procedure, results in an effective parameter-free method.

Finally, we remark that one could be tempted to use this idea, "as is", to stabilize the *symmetric* Lanczos recurrence. Indeed, although in exact arithmetic the Lanczos iteration applied to a symmetric A generates an orthonormal basis, computations in finite precision arithmetic cause loss of orthogonality; this fact has been analyzed in great detail starting with the seminal PhD thesis work of C.C. Paige, see [50]. Some preliminary experiments, however, seem to indicate that our procedure is not appropriate, at least in its current general form. Indeed, the matrix $\widehat{B}_k$ determined in Proposition 6.2 is nonsymmetric also in the case when B_k is symmetric, and complex eigenvalues do arise during the recurrence, irrespective of the possible loss of orthogonality in the Lanczos basis. Nonetheless, we believe that the topic is of great interest, and it grants a more specialized analysis.

Acknowledgments

We thank two anonymous reviewers for their insightful remarks. The work of V.S. was partially funded by the European Union—NextGenerationEU under the National Recovery and Resilience Plan (PNRR)—Mission 4 Education and research—Component 2 From research to business—Investment 1.1 Notice Prin 2022—DD N. 104 of 2/2/2022, entitled "Low-rank Structures and Numerical Methods in Matrix and Tensor Computations and their Application", code 20227PCCKZ—CUP J53D23003620006. V.S. is a member of the INdAM Research Group GNCS; its continuous support is gladly acknowledged. The work of Y.W. is funded by the National Natural Science Foundation of China (NSFC) Key Project (Grant No. 12131004); this work was carried out during a one-year visit of Y.W. at the University of Bologna, granted by a China Scholarship Council (File No. 202306860051). Open access publishing facilitated by Universita degli Studi di Bologna, as part of the Wiley - CRUI-CARE agreement.

Data Availability Statement

The authors have nothing to report.

Endnotes

[1] Bold face is used to highlight the fact that all basis vectors are orthonormal.

[2] A coupled two-term recurrence may be preferred for stability purposes, see, for example, [33].

[3] See, for instance, the right plot of Figure 6.1 in Reference [40], where a similar discussion was performed in terms of eigenvectors.

[4] The notation $\|S(Ux)\|_2^2 = (1 \pm \varepsilon)\|Ux\|_2^2$ is equivalent to $(1-\varepsilon)\|Ux\|_2^2 \leq \|S(Ux)\|_2^2 \leq (1+\varepsilon)\|Ux\|_2^2$, see [42].

[5] In matlab, `prob=rand(s,1); prob=prob/norm(prob,1).`

[6] Available at https://personalpages.manchester.ac.uk/staff/david.silvester/ifiss/.

[7] It is sufficient to compute the singular values; in MATLAB it reads `svd(orth(`$\widehat{U}_k$`'`$\boldsymbol{U}_k$`))` [10, section 6.4.3].

References

1. J. Liesen and Z. Strakos, *Krylov Subspace Methods. Principles and Analysis* (Oxford University Press, 2013).

2. Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd ed. (SIAM, Society for Industrial and Applied Mathematics, 2003).

3. M. Eiermann and O. Ernst, "A Restarted Krylov Subspace Method for the Evaluation of Matrix Functions," *SIAM Journal on Numerical Analysis* 44 (2006): 2481–2504.

4. A. Frommer, S. Güttel, and M. Schweitzer, "Efficient and Stable Arnoldi Restarts for Matrix Functions Based on Quadrature," *SIAM Journal on Matrix Analysis and Applications* 35 (2014): 661–683.

5. Y. Saad, *Numerical Methods for Large Eigenvalue Problems* (Society for Industrial and Applied Mathematics, 2011).

6. V. Simoncini and D. B. Szyld, "Recent Computational Developments in Krylov Subspace Methods for Linear Systems," *Numerical Linear Algebra With Applications* 14 (2007): 1–59.

7. D. C. Sorensen and C. Yang, "A Truncated QR Iteration for Large Scale Eigenvalue Calculations," *SIAM Journal on Matrix Analysis and Applications* 19 (1998): 1045–1073.

8. D. S. Watkins, *The Matrix Eigenvalue Problem. GR and Krylov Subspace Methods* (SIAM, 2007).

9. V. Simoncini and D. B. Szyld, "The Effect of Non-Optimal Bases on the Convergence of Krylov Subspace Methods," *Numerische Mathematik* 100 (2005): 711–733.

10. G. Golub and C. F. Van Loan, *Matrix Computations*, 4th ed. (Johns Hopkins University Press, 2013).

11. M. H. Gutknecht, "A Completed Theory of the Unsymmetric Lanczos Process and Related Algorithms. *I*," *SIAM Journal on Matrix Analysis and Applications* 13 (1992): 594–639.

12. O. Balabanov and L. Grigori, "Randomized Gram–Schmidt Process With Application to GMRES," *SIAM Journal on Scientific Computing* 44 (2022): A1450–A1474.

13. L. Burke, S. Güttel, and K. M. Soodhalter, "GMRES With Randomized Sketching and Deflated Restarting. Technical Report," 2023 arXiv n.2311.14206.

14. Y. Jang, L. Grigori, E. Martin, and C. Content, "Randomized Flexible GMRES With Deflated Restarting," *Numerical Algorithms* 98 (2025): 431–465.

15. Y. Nakatsukasa and J. A. Tropp, "Fast and Accurate Randomized Algorithms for Linear Systems and Eigenvalue Problems," *SIAM Journal on Matrix Analysis and Applications* 45 (2024): 1183–1214.

16. V. Rokhlin and M. Tygert, "A Fast Randomized Algorithm for Overdetermined Linear Least-Squares Regression," *Proceedings of the National Academy of Sciences* 105 (2008): 13212–13217.

17. O. Balabanov and A. Nouy, "Randomized Linear Algebra for Model Reduction. Part I: Galerkin Methods and Error Estimation," *Advances in Computational Mathematics* 45 (2019): 2969–3019.

18. A. Cortinovis, D. Kressner, and Y. Nakatsukasa, "Speeding up Krylov Subspace Methods for Computing f(A)b via Randomization," *SIAM Journal on Matrix Analysis and Applications* 45 (2024): 619–633.

19. S. Güttel and M. Schweitzer, "Randomized Sketching for Krylov Approximations of Large-Scale Matrix Functions," *SIAM Journal on Matrix Analysis and Applications* 44 (2023): 1073–1095.

20. D. Palitta, M. Schweitzer, and V. Simoncini, "Sketched and Truncated Polynomial Krylov Methods: Matrix Sylvester Equations. Technical Report 2311.16019, ArXiv," 2023 To Appear in Math. Comp.

21. G. W. Stewart, "An Arnoldi-Schur Algorithm for Large Eigenproblems," *SIAM Journal on Numerical Analysis* 23 (2001): 601–614.

22. V. Simoncini, "On the Convergence of Restarted Krylov Subspace Methods," *SIAM Journal on Matrix Analysis and Applications* 22 (2000): 430–452.

23. B. Philippe and L. Reichel, "On the Generation of Krylov Subspace Bases," *Applied Numerical Mathematics* 62, no. 9 (2012): 1171–1186Numerical Analysis and Scientific Computation With Applications (NASCA)., https://doi.org/10.1016/j.apnum.2010.12.009.

24. E. C. Carson, "Communication-Avoiding Krylov Subspace Methods in Theory and Practice, PhD Thesis, UC Berkeley," (2015).

25. S. Güttel and M. Schweitzer, "A Comparison of Limited-Memory Krylov Methods for Stieltjes Functions of Hermitian Matrices," *SIAM Journal on Matrix Analysis and Applications* 42 (2021): 83–107.

26. S. L. Campbell and C. D. Meyer, *Generalized Inverses of Linear Transformations*, Dover ed. (Dover, 1991).

27. Z. Bai, D. Day, and Q. Ye, "ABLE: An Adaptive Block Lanczos Method for Non-Hermitian Eigenvalue Problems," *SIAM Journal on Matrix Analysis and Applications* 20 (1999): 1060–1082.

28. C. Brezinski and H. Sadok, "Lanczos-Type Algorithms for Solving Systems of Linear Equations," *Applied Numerical Mathematics* 11 (1993): 443–473.

29. R. W. Freund, G. H. Golub, and N. M. Nachtigal, "Iterative Solution of Linear Systems," *Acta Numerica* 1 (1992): 57–100.

30. A. C. Antoulas, "Approximation of Large-Scale Dynamical Systems," in *Advances in Design and Control* (SIAM, 2005).

31. R. W. Freund, "Model Reduction Methods Based on Krylov Subspaces," *Acta Numerica* 12 (2003): 126–132.

32. E. Grimme, "Krylov Projection Methods for Model Reduction, PhD Thesis, the University of Illinois at Urbana-Champaign," 1997.

33. R. W. Freund and N. M. Nachtigal, "An Implementation of the QMR Method Based on Coupled Two-Term Recurrences," *SIAM Journal on Scientific Computing* 15 (1994): 313–337.

34. D. L. Boley and G. H. Golub, "The Nonsymmetric Lanczos Algorithm and Controllability," *Systems & Control Letters* 16 (1991): 97–105.

35. D. Day, "An Efficient Implementation of the Nonsymmetric Lanczos Algorithm," *SIAM Journal on Matrix Analysis and Applications* 18 (1997): 566–589.

36. C. Brezinski, M. Redivo Zaglia, and H. Sadok, "A Breakdown-Free Lanczos Type Algorithm for Solving Linear Systems," *Numerische Mathematik* 63 (1992): 29–38.

37. R. Freund, M. Gutknecht, and N. Nachtigal, "An Implementation of the Look-Ahead Lanczos Algorithm for Non-Hermitian Matrices," *SIAM Journal on Scientific Computing* 14 (1993): 137–158.

38. B. N. Parlett, D. R. Taylor, and Z. A. Liu, "A Look-Ahead Lanczos Algorithm for Unsymmetric Matrices," *Mathematics of Computation* 44 (1985): 105–124.

39. M. Crouzeix and C. Palencia, "The Numerical Range Is a $(1+\sqrt{2})$-Spectral Set," *SIAM Journal on Matrix Analysis and Applications* 38 (2017): 649–655.

40. D. Palitta, M. Schweitzer, and V. Simoncini, "Sketched and Truncated Polynomial Krylov Methods: Evaluation of Matrix Functions," *Numerical Linear Algebra With Applications* 32 (2025): e2596.

41. P.-G. Martinsson and J. A. Tropp, "Randomized Numerical Linear Algebra: Foundations and Algorithms," *Acta Numerica* 29 (2020): 403–572.

42. D. P. Woodruff, "Sketching as a Tool for Numerical Linear Algebra," *Foundations and Trends® in Theoretical Computer Science* 10 (2014): 1–157.

43. T. Sarlós, "Improved Approximation Algorithms for Large Matrices via Random Projections," in *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)* (IEEE, 2006), 143–152.

44. G. W. Stewart and J.-G. Sun, *Matrix Perturbation Theory* (Academic Press, 1990).

45. M. Eiermann and O. G. Ernst, "Geometric Aspects of the Theory of Krylov Subspace Methods," *Acta Numerica* 10 (2001): 251–312.

46. D. Kressner, "Memory-Efficient Krylov Subspace Techniques for Solving Large-Scale Lyapunov Equations," in *2008 IEEE International Conference on Computer-Aided Control Systems, San Antonio, TX, USA* (IEEE, 2008), 613–618.

47. R. A. Horn and C. R. Johnson, *Topics in Matrix Analysis* (Cambridge University Press, 1991).

48. H. C. Elman, D. J. Silvester, and A. J. Wathen, "Finite Elements and Fast Iterative Solvers, With Applications in Incompressible Fluid Dynamics," in *Numerical Mathematics and Scientific Computation*, vol. 21 of, ii ed. (Oxford University Press, 2014).

49. H. C. Elman, A. Ramage, and D. J. Silvester, "IFISS: A Matlab Toolbox for Modelling Incompressible Flow," *ACM Transactions on Mathematical Software* 33 (2007): 14.

50. C. Paige, "The Computation of Eigenvalues and Eigenvectors of Very Large Sparse Matrices, PhD Thesis, University of London," 1971.