

RESEARCH ARTICLE

Sketched and Truncated Polynomial Krylov Methods: Evaluation of Matrix Functions

Davide Palitta¹  | Marcel Schweitzer² | Valeria Simoncini^{1,3}

¹School of Mathematics and Natural Sciences, Bergische Universität Wuppertal, Wuppertal, Germany | ²Dipartimento di Matematica, (AM)², Alma Mater Studiorum - Università di Bologna, Bologna, Italy | ³IMATI-CNR, Pavia, Italy

Correspondence: Davide Palitta (davide.palitta@unibo.it)

Received: 28 March 2024 | **Revised:** 24 September 2024 | **Accepted:** 30 September 2024

Funding: This work was supported by Gruppo Nazionale per il Calcolo Scientifico (project GNCS2022 “Tecniche Avanzate per Problemi Evolutivi: Discretizzazione, Algebra Lineare Numerica, Ottimizzazione”) (CUP_E55F22000270001).

Keywords: Krylov subspace method | matrix function | projection method | randomization | sketching | subspace embedding

ABSTRACT

Among randomized numerical linear algebra strategies, so-called sketching procedures are emerging as effective reduction means to accelerate the computation of Krylov subspace methods for, for example, the solution of linear systems, eigenvalue computations, and the approximation of matrix functions. While there is plenty of experimental evidence showing that sketched Krylov solvers may dramatically improve performance over standard Krylov methods, especially when combined with a truncated orthogonalization step, many features of these schemes are still unexplored. We derive a new sketched Arnoldi-type relation that allows us to obtain several different new theoretical results. These lead to an improvement of our understanding of sketched Krylov methods, in particular by explaining why the frequently occurring sketched Ritz values far outside the spectral region of A do not negatively influence the convergence of sketched Krylov methods for $f(A)b$. Our findings also help to identify, among several possible equivalent formulations, the most suitable sketched approximations according to their numerical stability properties. These results are also employed to analyze the error of sketched Krylov methods in the approximation of the action of matrix functions, significantly contributing to the theory available in the current literature.

1 | Introduction

The field of numerical linear algebra is currently experiencing a shift of paradigm with the advent of fast and scalable randomized algorithms for many core linear algebra problems, the most prominent one probably being the now widely adopted randomized singular value decomposition [1]. A specific, important tool that has emerged in recent years in the context of randomized numerical linear algebra is the so-called *sketch-and-solve* paradigm. While initially only used—very successfully—for computations with rather crude accuracy demands, for example, (low-rank) matrix approximation [2, 3], it was recently discovered that—when combined with Krylov subspace methods—it might

also be applicable for high accuracy computations as, for example, the solution of linear systems, eigenvalue computations, and the approximation of matrix functions [4–7]. The main idea in these methods is to use randomized sketching for reducing the cost of orthogonalization by, for example, combining it with a truncated Gram–Schmidt process. This often allows a huge boost in performance, as orthogonalization is typically the main bottleneck in polynomial Krylov methods when applied to non-symmetric problems, particularly on modern high-performance architectures, where each inner product introduces a global synchronization point and communication tends to be the dominating factor. For symmetric matrices, computation of an orthogonal Krylov basis is possible by the short recurrence Lanczos method

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2024 The Author(s). *Numerical Linear Algebra with Applications* published by John Wiley & Sons Ltd.

[8], so that the gains one can obtain by reducing the cost of orthogonalization are rather limited. We therefore focus on the nonsymmetric case in the following.

While it was experimentally confirmed that sketched Krylov solvers might dramatically improve performance over standard Krylov methods [4–7], many features of these methods are not fully understood yet. In particular, the numerical stability of these methods and how exactly they relate to their non-randomized counterparts needs further analysis. Several recent studies lay a theoretical groundwork with elements of an analysis [5, 6, 9, 10], however, none of these can satisfactorily explain the behavior observed in practice yet, in particular for nonsymmetric problems. For example, sometimes these methods work incredibly well, essentially reproducing the behavior of the full (non-sketched) Arnoldi method at a fraction of the cost, but in other cases they may fail completely; see, for example, the experiments reported in [5, 6]. The currently available theory can neither satisfactorily explain the successful cases, nor the failures.

In this work, we derive new theoretical results which help better understand the overall behavior of Krylov methods that combine truncated orthogonalization with randomized sketching and identify which of the several possible (mathematically equivalent) formulations of sketched Krylov approximations can be expected to enjoy the most favorable numerical stability properties. We then use these results to analyze the error of sketched Krylov methods for approximating the action of matrix functions, $f(A)b$, where $A \in \mathbb{R}^{n \times n}$ and $b \in \mathbb{R}^n$, significantly contributing to the currently available theory [5, 6].

One of our main results is a new Arnoldi-type relation for the sketched basis (Proposition 3.1). This is a general result and it can be used to analyze many of the sketching approaches obtained so far, from linear systems to eigenvalue problems. We show that full, truncated, and sketched Arnoldi methods are all related by basis transformations and rank-one additions; see the summary in Table 1. These explicit relations allow us to derive new algebraic expressions for the distance between the full Arnoldi and sketched approximations to matrix function evaluation, which we analyze in detail (Corollaries 5.1 and 5.4). This analysis is based on new results concerning the connection between functions of certain rank-one modified matrices and divided differences (Theorems 4.1 and 4.2). We also provide an explanation of the good performance of the so-called sketched-whitened Arnoldi method, in spite of apparently adverse spectral properties of the involved projected matrix; see in particular the family of bounds (42).

Here is an outline of the paper. In Section 2 we recall some background material that will serve as a foundation for our analysis. In particular, we cover the (full and truncated) Arnoldi method for $f(A)b$ (Section 2.1), subspace embeddings (Section 2.2), and

prior work on sketched Krylov methods (Section 2.3). In Section 3 we show that a new Arnoldi-like relation, that we name the *sketched Arnoldi relation*, holds also when truncated Krylov subspace methods and sketching techniques are combined. Then, in Section 4 we show how to rewrite the sketched Arnoldi relation in terms of functions of certain rank-one modified matrices and present some general results on such matrix functions that will be employed to relate sketched and non-sketched approximations in Section 5. The role of eigenvector-related coefficients is further investigated in Section 6, leading to the counterintuitive observation that the outlying spurious Ritz values generated by the sketched approach do not contribute to the error vector. In Section 7 we discuss some issues related to the stability of the analyzed procedures, while our conclusions are given in Section 8.

In the following, the notation $\| \cdot \|$ will be used to denote the Euclidean norm for vectors and the corresponding induced norm for matrices. Similarly, $\| z \|_\infty := \max_i |z_i|$ will denote the vector infinity norm whereas $\| T \|_\infty$ the corresponding matrix norm. The identity matrix will be denoted by I , with dimension clear from the context, while e_j will denote the j th column of the identity matrix. We denote the d th Krylov subspace corresponding to $A \in \mathbb{R}^{n \times n}$ and $b \in \mathbb{R}^n$ by $\mathcal{K}_d(A, b) = \text{span}\{b, Ab, A^2b, \dots, A^{d-1}b\}$. Finally, roman font ($\mathcal{H}_d, \mathcal{U}_d$) will be used for quantities related to a non-orthogonal Arnoldi relation and calligraphic font ($\underline{\mathcal{H}}_d, \underline{\mathcal{U}}_d$) for quantities related to a fully orthogonal Arnoldi relation. Matlab [11] has been used for all our numerical experiments.

Throughout, exact arithmetic computations will be assumed in our derivations.

2 | Background Material

In this section, we review a few known tools that will be crucial ingredients in our analysis.

2.1 | The Full and Truncated Arnoldi Method

The backbone of most Krylov subspace methods is the Arnoldi process [12], which, given a matrix $A \in \mathbb{R}^{n \times n}$, a vector $b \in \mathbb{R}^n$, and a target dimension d , computes a nested orthonormal basis $\mathcal{U}_d \in \mathbb{R}^{n \times d}$ of $\mathcal{K}_d(A, b)$ via the modified Gram–Schmidt orthonormalization and gives rise to the *Arnoldi relation*

$$A\mathcal{U}_d = \mathcal{U}_{d+1}\underline{\mathcal{H}}_d = \mathcal{U}_d\mathcal{H}_d + h_{d+1,d}u_{d+1}e_d^T \quad (1)$$

where $\underline{\mathcal{H}}_d = [\mathcal{H}_d; h_{d+1,d}e_d^T] \in \mathbb{R}^{(d+1) \times d}$, is an unreduced upper Hessenberg matrix containing the orthogonalization coefficients, and $\mathcal{U}_{d+1} = [\mathcal{U}_d, u_{d+1}]$.

TABLE 1 | Summary of analyzed matrix function approximations, in the truncated and fully orthonormal bases.

Approx	Truncated basis U_d	Fully orthonormal basis \mathcal{U}_d
f_d^{FOM}	—	$\mathcal{U}_d f(\mathcal{H}_d) e_1 \ b \ $ (2)
f_d^{TR}	$U_d f(H_d) e_1 \ b \ $ (4)	$\mathcal{U}_d f(\mathcal{H}_d - h_{d+1,d} t e_d^T) e_1 \ b \ $ (36)
f_d^{SK}	$U_d T_d^{-1} f(\hat{H}_d + \hat{t} e_d^T) e_1 \ S b \ $ (18)	$\mathcal{U}_d f(\mathcal{H}_d + \hat{r} e_d^T) e_1 \ b \ $ (33)

Given the quantities from the Arnoldi relation (1), one can extract an approximation for $f(A)b$ from the Krylov subspace via

$$f_d^{\text{FOM}} = \mathcal{U}_d f(\mathcal{K}_d) e_1 \| b \| \quad (2)$$

We call this the full Arnoldi (or FOM, in analogy with the Full Orthogonalization Method for linear systems) approximation.

The main contributions to the computational cost of performing d steps of the Arnoldi method are d matrix vector products with A at a cost of $\mathcal{O}(d \text{ nnz}(A))$ floating point operations (flops), assuming that A is sparse with $\text{nnz}(A)$ nonzeros, and the modified Gram–Schmidt orthogonalization which has a cost of $\mathcal{O}(d^2 n)$ flops across all d iterations. Due to its quadratic growth, the cost of orthogonalization quickly dominates the overall cost of the method. In particular, on modern high performance architectures, each inner product introduces a global synchronization point, which leads to a large amount of communication, unless specific “low-sync” implementations are used, which come with their own challenges regarding numerical stability; see, for example, the recent survey [13].

An alternative—which is well-explored in the solution of linear systems—is to use a truncated Arnoldi method (see algorithm 6.6 in [14]), that is, instead of orthogonalizing a new basis vector against all previous basis vectors, one only orthogonalizes it against a fixed number k of previous vectors (with $k = 2$ mimicking the Lanczos method for Hermitian A [8]), thereby overcoming the quadratic growth of the orthogonalization cost. This leads to a *truncated Arnoldi relation*

$$AU_d = U_{d+1} \underline{H}_d \quad (3)$$

with $U_{d+1} = [U_d, u_{d+1}]$, where now \underline{H}_d is banded and the columns of U_d are not all mutually orthonormal. In the following we denote $\underline{H}_d = [H_d; h_{d+1,d} e_d^T]$, where only the last component in the last row is nonzero. Given (3), one can compute the truncated FOM approximation

$$f_d^{\text{TR}} = U_d f(H_d) e_1 \| b \| \quad (4)$$

which was, for example, recently proposed in algorithm 3.3 of [5]. The price one pays for the reduction in orthogonalization cost when using (4) is that, typically, convergence is delayed in comparison to the full Arnoldi approximation (2); see, for example [5] as well as Example 7.1 below.

2.2 | Subspace Embeddings

The foundations of the sketched Krylov methods discussed in this paper are *subspace embeddings* [2, 15, 16], which embed a low-dimensional subspace \mathcal{V} of \mathbb{R}^n into a smaller space \mathbb{R}^s , $s \ll n$, such that inner products (or norms) are distorted in a controlled manner. Specifically, a matrix $S \in \mathbb{R}^{s \times n}$ is called an ε -subspace embedding for \mathcal{V} (for a given distortion parameter $\varepsilon \in [0, 1)$), if

$$(1 - \varepsilon) \| v \|^2 \leq \| Sv \|^2 \leq (1 + \varepsilon) \| v \|^2 \quad (5)$$

for all $v \in \mathcal{V}$. This is equivalent to requiring

$$\langle u, v \rangle - \varepsilon \| u \| \| v \| \leq \langle Su, Sv \rangle \leq \langle u, v \rangle + \varepsilon \| u \| \| v \| \quad (6)$$

for all $u, v \in \mathcal{V}$ by the parallelogram inequality.

Constructing S as above requires knowledge of the subspace \mathcal{V} that shall be embedded. In our context, \mathcal{V} will be a Krylov subspace $\mathcal{K}_d(A, b)$ which is not known in advance. To still construct an embedding in this case, one can use so-called *oblivious embeddings*. A matrix S is an oblivious ε -subspace embedding for subspaces of dimension d , if condition (5) holds with high probability for *any* subspace \mathcal{V} with $\dim(\mathcal{V}) = d$. To be precise, the inputs for constructing S are the embedding quality ε , the sketching dimension s and an accepted failure probability δ . Based on these, S is constructed using probabilistic methods; see, for example [16], or sect. 8 of [17] for details.

Frequently, oblivious subspace embeddings are constructed as subsampled trigonometric transforms,

$$S = \sqrt{\frac{s}{n}} DNE \quad (7)$$

where $E \in \mathbb{R}^{n \times n}$ is a diagonal matrix with Rademacher entries (i.e., the diagonal entries are randomly chosen as ± 1 with equal probability), $D \in \mathbb{R}^{s \times n}$ contains s randomly selected rows of the identity matrix, and $N \in \mathbb{R}^{n \times n}$ is an orthogonal trigonometric transform (e.g., discrete Fourier transform, discrete cosine transform or Walsh–Hadamard transform). One can show that (7) with $s = \mathcal{O}\left(\varepsilon^{-2} \left(d + \log \frac{n}{\delta}\right) \log \frac{d}{\delta}\right)$ is an oblivious ε -subspace embedding with failure probability δ . With a careful implementation, such S can be applied to a vector in $\mathcal{O}(n \log s)$ flops. In practice, selecting the smaller sketching dimension $s = \mathcal{O}\left(\varepsilon^{-2} \frac{d}{\delta}\right)$ appears to work very well, except for a few academic examples. We refer to, for example [18], and sect. 9 of [17] for a discussion and experimental evidence.

As the required sketching dimension s depends quadratically on the embedding quality ε , one typically needs to accept a rather crude embedding quality, such as $\varepsilon = 1/\sqrt{2}$ or $\varepsilon = 1/2$, in order to avoid excessive growth of computational cost.

To conclude this subsection, we mention that a subspace embedding S induces a semidefinite inner product

$$\langle u, v \rangle_S := \langle Su, Sv \rangle \quad (8)$$

From the embedding property, it directly follows that when restricted to \mathcal{V} , $\langle \cdot, \cdot \rangle_S$ is an actual inner product; see, for example, sect. 3.1 in [9].

In the remainder of the paper, we often simply call S a “sketching matrix,” thereby meaning that it is an ε -subspace embedding for a certain Krylov subspace (which will be clear from the context) with a suitable value of ε .

2.3 | The Sketched Arnoldi Method

As already mentioned in the introduction, a rather novel approach for reducing the orthogonalization cost in Krylov subspace methods via truncated orthogonalization *without* sacrificing convergence speed is to combine the method with oblivious subspace embeddings, leading to so-called *sketched Krylov methods*; see [7] for an application to linear systems and eigenvalue problems as well as [5, 6] for approximating the action of matrix

functions. Similar ideas have also been discussed in [9] in the context of model order reduction.

The main idea of the methods in [6, 7] is to first cheaply construct a non-orthogonal Krylov basis and then subsequently approximately orthogonalize it, working just with sketched vectors of size $s \ll n$ (a process known as *basis whitening*). We explain the methodology for approximations of $f(A)b$, closely following sect. 2 of [6]. Given a non-orthogonal Krylov basis U_d resulting from a truncated Arnoldi process and a sketching matrix S , the “sketched FOM” (sFOM) approximation for $f(A)b$ can be defined as

$$f_d^{\text{SK}} = U_d f \left((SU_d)^\dagger SAU_d \right) (SU_d)^\dagger Sb \quad (9)$$

where $(SU_d)^\dagger$ denotes the Moore-Penrose pseudoinverse of SU_d . The basis whitening procedure then consists of computing a thin QR decomposition $SU_d = Q_d T_d$ of the sketched basis and performing the replacements

$$SU_d \leftarrow Q_d, \quad SAU_d \leftarrow (SAU_d)T_d^{-1}, \quad U_d \leftarrow U_d T_d^{-1} \quad (10)$$

Note that the matrix $U_d T_d^{-1}$ from the last replacement need not be formed,¹ because inserting (10) into (9) yields

$$f_d^{\text{SK}} = U_d (T_d^{-1} f(Q_d^* SAU_d T_d^{-1}) Q_d^* S b) \quad (11)$$

which can be evaluated by first applying T_d^{-1} via back substitution and then multiplying by U_d .

The motivation for using (10) and (11) is that the transformed basis $U_d T_d^{-1}$ appearing in (16) is orthogonal with respect to the inner product $\langle \cdot, \cdot \rangle_S$ and is known to be extremely well-conditioned in terms of its spectral condition number. This is stated in the following result, which directly follows from Corollary 2.2 in [4]; see also sect. 4 in [19] for related, earlier results in a probabilistic setting.

Proposition 2.1. *If S is an ε -subspace embedding of $\mathcal{X}_d(A, b)$, then the 2-norm condition number of $U_d T_d^{-1}$ is bounded as*

$$\kappa_2(U_d T_d^{-1}) \leq \sqrt{\frac{1+\varepsilon}{1-\varepsilon}} \quad (12)$$

For example, for a practically reasonable distortion parameter $\varepsilon = 1/\sqrt{2}$, Proposition 2.1 guarantees that $\kappa_2(U_d T_d^{-1}) \leq 1 + \sqrt{2} \approx 2.4142$.

Numerical experiments in sect. 5 of [6] show that indeed, the approximation (11) often converges remarkably similar to the full Arnoldi approximation despite using the cheaper truncated orthogonalization. It is also observed that the method may sometimes fail completely, though.

As an alternative to using sketching as a means to retrospectively improve the conditioning of a locally orthogonal Krylov basis, another recently explored avenue for reducing orthogonalization cost and communication in Krylov methods via sketching is to perform a *full* Gram–Schmidt orthogonalization, but with respect to the semidefinite inner product $\langle \cdot, \cdot \rangle_S$ defined in

(8) instead of the Euclidean inner product; cf., for example [4, 10], for use of this technique in the context of solving linear systems by FOM or GMRES. In algorithm 3.1 of [5] this approach is extended to approximating the action of matrix functions. By using this approach, inner products in the orthogonalization become cheaper (as they involve only vectors of length $s \ll n$), but the asymptotic cost of the orthogonalization remains at $\mathcal{O}(d^2 n)$, just as for full Arnoldi. This approach is therefore mainly attractive when working in a highly parallel computing environment where inner products are the main bottleneck due to expensive communication.

3 | The Sketched Arnoldi Relation

In the following, we focus on better understanding the behavior of the sketched Arnoldi approximation (9) (or equivalently (11)). To do so, we start by deriving a sketched analogue of the Arnoldi relation (1), which allows to rewrite (11) in a way that is both easier to interpret and—more importantly—allows to derive several new theoretical results in later sections of this paper.

Proposition 3.1. *Let U_{d+1}, \underline{H}_d be the quantities from a truncated Arnoldi relation (3) and let $SU_{d+1} = Q_{d+1} T_{d+1}$ be a thin QR decomposition with*

$$Q_{d+1} = [Q_d, q] \quad \text{and} \quad T_{d+1} = \begin{bmatrix} T_d & t \\ 0^T & \tau_{d+1} \end{bmatrix}$$

Then the following Arnoldi-like formula holds:

$$SAU_d = SU_d (H_d + r e_d^T) + \tau_{d+1} h_{d+1,d} q e_d^T \quad (13)$$

with $r = h_{d+1,d} T_d^{-1} t$ and $q \perp SU_d$.

Proof. Let $h^T = h_{d+1,d} e_d^T$. We have

$$\begin{aligned} SAU_d &= SU_{d+1} \underline{H}_d = Q_{d+1} T_{d+1} \underline{H}_d = Q_{d+1} \begin{bmatrix} T_d H_d + t h^T \\ [0^T, \tau_{d+1}] \underline{H}_d \end{bmatrix} \\ &= Q_d T_d (H_d + T_d^{-1} t h^T) + \tau_{d+1} q h^T = SU_d (H_d + T_d^{-1} t h^T) + \tau_{d+1} q h^T \end{aligned} \quad (14)$$

The result follows by inserting the definition of r .

The action of sketching appears only in the vector r and in the scalar τ_{d+1} , while the matrix $H_d + r e_d^T$ differs from H_d only in its last column, with a rank-one modification. This will have a quite interesting impact in the following derivations.

Remark 1. Relations of the form (3) do not only arise in the context of the truncated Arnoldi method, but also in case of other recurrences for generating non-orthogonal Krylov bases, for example, employing Chebyshev or Newton polynomials; see [20]. Proposition 3.1 thus also applies to these methods, but we restrict to truncated Arnoldi here for ease of exposition and because it appears to be the most widely used basis generation method in sketched Krylov methods.

In order to relate Proposition 3.1 to the basis whitening procedure (10) outlined in Section 2.3, first note that this procedure

was already implicitly employed in the derivation of the sketched Arnoldi relation due to occurrence of the thin QR decomposition. The underlying orthogonalization step can be even better appreciated by rewriting (14) as

$$SAU_d T_d^{-1} = Q_d T_d (H_d + T_d^{-1} t h^T) T_d^{-1} + \tau_{d+1} q h^T T_d^{-1} \quad (15)$$

from which we obtain the *whitened-sketched* Arnoldi relation (WS-Arnoldi in short)

$$SA(U_d T_d^{-1}) = Q_d (\hat{H}_d + \hat{t} e_d^T) + \frac{h_{d+1,d} \tau_{d+1}}{\tau_d} q e_d^T \quad (16)$$

where $\hat{H}_d = T_d H_d T_d^{-1}$ and $\hat{t} = (h_{d+1,d} / \tau_d) t$.

The inclusion of the orthogonalization coefficients T_d , and thus the use of (16) in place of (13), seems to be very beneficial to the numerical computation. Indeed, computing functions of $H_d + r e_d^T$ appears to be quite sensitive to round-off, due to the growing components of r , which lead to poor balancing; see the discussion in Remark 3. Note however, that

$$\sqrt{\frac{1-\varepsilon}{1+\varepsilon}} \kappa_2(T_d) \leq \kappa_2(U_d) \leq \sqrt{\frac{1+\varepsilon}{1-\varepsilon}} \kappa_2(T_d)$$

so that T_d must necessarily be ill-conditioned whenever U_d is ill-conditioned. As U_d results from a truncated orthogonalization process, it will become more and more ill-conditioned as d increases. Despite this fact, inversion of T_d can be expected to not be very problematic in the context of sketched Krylov methods, as we discuss in Section 7.

The result of Proposition 3.1 can be used to analyze sketched Krylov methods for the approximation of matrix functions $f(A)b$. Some first steps regarding the analysis of these methods were done in sect. 3.2 of [6] and sect. 4 of [5]. Still, sketched Krylov methods for $f(A)b$ are far from being fully understood from a theoretical perspective. We highlight two particular issues that have not been well explored so far. First, sketching might (and often does) introduce “sketched Ritz values” which lie very far outside the spectral region of A . However, these tend to mostly not negatively influence the performance of the methods. Second, there exist many different formulations for how to compute sketched Krylov approximations, all of them involving computations with potentially very ill-conditioned matrices; cf. also the discussion in the preceding paragraph. Still, even in the presence of severe ill-conditioning and rank deficiency, sketched Krylov methods can work well in practice. Throughout the next sections, we approach both these issues and derive new theoretical results which allow to better capture the properties of the sketched Arnoldi approximation and better explain the observed behavior. We start by discussing the relation between the sketched Arnoldi approximation and functions of certain rank-one modified matrices in Section 4.

4 | The Impact of Rank-One Modifications on Matrix Functions

The formula (11) defining the sketched Arnoldi approximation is quite hard to interpret and thus not very insightful. In particular, the meaning of the small matrix $Q_d^* S A U_d T_d^{-1}$ at which f is

evaluated is rather unclear. We now show how to rewrite it using the tools from Section 3 to obtain formulas that are both easier to interpret and more suitable for a subsequent theoretical analysis.

From the sketched Arnoldi relation (13) and the fact that $(S U_d)^\dagger = T_d^{-1} Q_d^*$, we directly obtain by straightforward algebraic manipulations that (9) simplifies to

$$f_d^{\text{SK}} = U_d f(H_d + r e_d^T) e_1 \| b \| \quad (17)$$

a form much more reminiscent of the standard and truncated FOM approximations (2) and (4). Using the WS-Arnoldi relation (16), we can also write

$$f_d^{\text{SK}} = U_d T_d^{-1} f(\hat{H}_d + \hat{t} e_d^T) e_1 \| S b \| \quad (18)$$

As shown above, the sFOM approximation to $f(A)b$ involves a rank-one modification to H_d , or a similar matrix. To better understand the role of such modification, we begin by stating a quite general result, which provides an explicit relationship between $f(M) \in \mathbb{R}^{d \times d}$ and functions of a rank-one modification of the form $f(M + v e_d^T)$. We assume in the following that $M + v e_d^T$ is diagonalizable with eigenvalue decomposition

$$M + v e_d^T = X \Lambda X^{-1}, \quad \text{where } \Lambda = \text{diag}(\lambda_1, \dots, \lambda_d) \quad (19)$$

and define the corresponding function

$$g_v(t) = \sum_{i=1}^d \alpha_i \beta_i f[t, \lambda_i] \quad (20)$$

where

$$f[t, \lambda] = \begin{cases} f'(\lambda), & \text{if } t = \lambda \\ \frac{f(t) - f(\lambda)}{t - \lambda}, & \text{otherwise} \end{cases} \quad (21)$$

denotes a divided difference and $\alpha_i = e_d^T X e_i$, $\beta_i = e_i^T X^{-1} e_1$ with X from the eigenvalue decomposition (19) of $M + v e_d^T$. The subscript in g_v is meant to keep track of the rank-one modification used to generate the spectral quantities defining g_v . The proof of this result is given in Appendix B.

Theorem 4.1. *Let $M + v e_d^T$ have the eigendecomposition (19) and assume that f is defined in a region containing the eigenvalues of M and $M + v e_d^T$. Then*

$$f(M + v e_d^T) e_1 - f(M) e_1 = g_v(M) v$$

where g_v is defined in (20).

The formula derived in Theorem 4.1 provides insightful information on the role of the (undesired) eigenvalues of the reduced sketched matrix; see Section 6. However, this expression for the matrix function g_v makes it difficult to analyze and measure the error between the different function evaluations; see also [21] where such observations are made for g_0 . The following reformulation leads to a more penetrating interpretation, and it may be read off as a different definition of g_v ; see Appendix B for its proof. Notice that the second definition in the theorem below is stated in terms of higher-order divided differences, which — among other

representations—can be defined recursively based on (21) (see, e.g., sect. 5 in [22]) or via the Cauchy integral formula, as

$$f[z_1, \dots, z_{d+1}] = \frac{1}{2\pi i} \int_{\Gamma} \frac{f(z)}{(z - z_1) \cdots (z - z_{d+1})} dz \quad (22)$$

see sect. 9 in [22].

Theorem 4.2. *Let M be upper Hessenberg and let $\mathbf{m}_d := \prod_{j=1}^{d-1} M_{j+1,j}$, where $M_{j+1,j}$ denotes the $(j+1, j)$ th entry of M . Assume that $M + ve_d^T$ has the eigendecomposition (19), that all eigenvalues $\{\lambda_1, \dots, \lambda_d\}$ of $M + ve_d^T$ are simple and that f is defined in a region containing the eigenvalues of M and $M + ve_d^T$. Then*

$$f(M + ve_d^T)e_1 - f(M)e_1 = \mathbf{m}_d \sum_{i=1}^d \frac{1}{\omega'(\lambda_i)} f[M, \lambda_i] v \quad (23)$$

with $\omega(z) = \prod_{i=1}^d (z - \lambda_i)$, from which it also holds

$$f(M + ve_d^T)e_1 - f(M)e_1 = \mathbf{m}_d f[M, \lambda_1, \dots, \lambda_d] v \quad (24)$$

where $f[M, \lambda_1, \dots, \lambda_d]$ is the order- $(d+1)$ divided difference of the function f .

Remark 2. Higher-order divided differences have been used before in matrix function computations, see, for example [23, 24], where they are employed for representing the error of the (standard) Arnoldi approximation in the context of restarted methods (where, in our notation, the function g_0 appears, as no rank-one modifications occur in that setting). As the authors of [23, 24] observe, the numerical computation of higher-order divided differences can quickly become unstable, so that restarted algorithms based on these error representations did not turn out to be universally applicable, in particular not for “hard” problems. Except for [21], where a slightly more stable representation based on first-order divided differences was used, similar to our Theorem 4.1, there has not been too much subsequent work exploring this connection. In sect. 4 of [24], it is also briefly discussed how to use a divided difference representation in order to estimate the norm of the Arnoldi error when A is (close to) normal.

While it might not be generally useful for numerical computations, in the following we rather use the divided difference representation as a theoretical tool which allows to compare different methods and obtain insight into their qualitative behavior.

A few comments on the previous results are in order. By comparing Theorem 4.1 and (23), it follows that for any $i \in \{1, \dots, d\}$ it holds

$$\alpha_i \beta_i = \frac{\prod_{j=1}^{d-1} M_{j+1,j}}{\omega'(\lambda_i)}$$

providing some insight into the magnitude of the product $\alpha_i \beta_i$ for outlying eigenvalues λ_i . In Section 6 the role of β_i will be explored in more details.

We now use the above analysis in order to obtain a general bound for the norm of the difference between a function of a matrix

and its rank-one modification. This bound emphasizes the role of the function f and the domain on which it acts. The role of the vector v defining the rank-one modification will be discussed in more detail in Section 5. The results in Theorem 4.2 can be conveniently interpreted by using a matrix analog of the well-known relation between divided differences and differentiation. For an analytic function h defined on the field of values of M , we will also make use of the following bound proved by Crouzeix and Palencia [25],

$$\|h(M)\|_2 \leq (1 + \sqrt{2}) \max_{z \in F(M)} |h(z)| \quad (25)$$

We recall that the field of values of a square matrix $M \in \mathbb{C}^{d \times d}$ is defined as $F(M) = \{z \in \mathbb{C} : z = (x^* M x) / (x^* x), x \in \mathbb{C}^d, x \neq 0\}$.

The following bound can be derived from the *Genocchi–Hermite representation* of divided differences; see, for example [22], appendix B.16 of [26] or problem (6.1.43) in [27].

Proposition 4.1. *Let D be a closed, convex set in the complex plane containing z_1, \dots, z_{d+1} and let f be analytic on D . Then*

$$|f[z_1, \dots, z_{d+1}]| \leq \frac{\max_{\zeta \in D} |f^{(d)}(\zeta)|}{d!}$$

From this proposition, we obtain the following bound which will be used in Section 6 below for the sketched Arnoldi approximation.

Corollary 4.1. *Let the assumptions of Theorem 4.2 hold and let $D \subseteq \mathbb{C}$ be convex and contain $F(M)$ as well as the nodes $\lambda_1, \dots, \lambda_d$. Then,*

$$\|f(M + ve_d^T)e_1 - f(M)e_1\| \leq \mathbf{m}_d (1 + \sqrt{2}) \frac{\max_{\zeta \in D} |f^{(d)}(\zeta)|}{d!} \|v\| \quad (26)$$

Proof. Denote

$$h = f[\cdot, \lambda_1, \dots, \lambda_d] \quad (27)$$

that is, $h(z) = f[z, \lambda_1, \dots, \lambda_d]$. From (24), we then have

$$\|f(M + ve_d^T)e_1 - f(M)e_1\| = \mathbf{m}_d \|h(M)\| \quad (28)$$

Due to the assumptions on D , we can apply Proposition 4.1 and further use (25) to bound

$$\|h(M)\| \leq (1 + \sqrt{2}) \max_{z \in F(M)} |h(z)| \leq (1 + \sqrt{2}) \frac{\max_{\zeta \in D} |f^{(d)}(\zeta)|}{d!} \quad (29)$$

Putting (28) and (29) together concludes the proof.

For bounded D , estimates of the form (26) are insightful for functions such as the exponential, for which it is easy to see that the term $\mathbf{m}_d / d!$ will eventually go to zero with growing d , as $\mathbf{m}_d \leq (\|M\|_\infty)^d$, while all derivatives $f^{(d)}$ remain bounded. Interestingly, this asymptotic behavior agrees with classical bounds for Krylov subspace approximations for the exponential function [28]. In the next section we will show that the error formulas in

Theorem 4.1 and Theorem 4.2 are very general, and can be used to relate the approximation obtained by the fully orthogonal basis with that computed by either the truncated or sketched bases. Hence, the bound in (26), although rather pessimistic in general, ensures that the two quantities $f(M + ve_d^T)e_1$ and $f(M)e_1$ will coincide in exact arithmetic, for large enough d .

Example 4.3. We consider the following problem with the Toeplitz matrix M_d , where the subscript stresses the dependence on the dimension d ,

```
for d=5:5:30
M=toeplitz([-4, 2, zeros(1,d-2)],
[-4, 1/2, 1/2, zeros(1,d-3)]);
rng(21*pi); e1=eye(d,1);
v=flipud(linspace(1,20,d))'.*randn(d,1);
Mhat=M+v*[zeros(1,d-1),1];
y1=expm(M)*e1;
y1s=expm(Mhat)*e1;
end
```

These data were chosen for full reproducibility. We also consider the case where the Toeplitz matrix M is substituted by the true upper Hessenberg matrix obtained with the full Arnoldi recurrence at the given iteration d . For this latter case, the Arnoldi iteration was used with b equal to a unit norm constant vector of dimension 2500, and A the finite difference discretization of the operator $\Delta u - 10xu_x - 10yu_y$, subject to zero boundary conditions. The left plot of Figure 1 shows the error norm $\| \exp(M_d + v_d e_d^T) e_1 - \exp(M_d) e_1 \|$ and the bound $(1 + \sqrt{2}) \| v_d \| \mathbf{m}_d / d!$. Since \mathbf{m}_d is moderate, the convergence slope of the asymptotic curve nicely follows that of the true error. The right plot shows an analogous performance for the matrix obtained via an Arnoldi procedure.

For functions other than the exponential the bound in (26) may not be significant. As an example, for $z^{-1/2}$ the formula for the derivative contains a factor that essentially grows as the factorial. To provide further insight, we resort to representations involving more general tools that take into account the properties of divided

differences in complex domains. More precisely, we can use a result proved by Curtiss in 1962. We first recall that for a bounded region D of the complex plane whose boundary is an admissible Jordan curve, it is possible to define a function $z = \chi(\omega) = \mu\omega + \mu_0 + \frac{\mu_1}{\omega} + \frac{\mu_2}{\omega^2} + \dots$, where $\mu > 0$ is called the capacity of the Jordan curve, and ω is defined outside the unit circle. The function χ maps conformally the region outside the unit circle into the exterior of \bar{D} . Setting $\omega = r \exp(i\theta)$, there exists a maximum value of r , that is $\rho > 1$, such that f is analytic on the region between the exterior of D and the curve $C_r = \{z = \chi(r \exp(i\theta))\}$, with $\theta \in [0, 2\pi]$. We refer to [29] for a more comprehensive description of the notation.

Theorem 4.4. (Theorem 4.1 in [29]) With the above notation, let D be a bounded region of the complex plane whose boundary is an admissible Jordan curve, and let μ be the capacity of ∂D . Let f be analytic on $C = \bar{D}$ and let the sequence of $\{\lambda_1, \dots, \lambda_d\}$ be equidistributed on ∂D . Let $\rho > 1$ be the largest value such that f is analytic in the extended curve C_ρ containing C . Then for any r , $1 < r < \rho$,

$$|f[\lambda_1, \dots, \lambda_d]| \leq \frac{\ell_{C_r}}{2\pi(\mu r)^d} \max_{z \in C_r} |f(z)|$$

where ℓ_{C_r} is the length of the curve C_r .

Combining the result of Theorem 4.4 with (24) and the first inequality in (29), we obtain the following bound, where D can be taken as the convex set above,

$$\begin{aligned} & \| f(M + ve_d^T) e_1 - f(M) e_1 \| \\ & \leq \mathbf{m}_d (1 + \sqrt{2}) \max_{z \in F(M)} |f[z, z_1, \dots, z_d]| \| v \| \\ & \sim \ell_{C_r} \frac{\mathbf{m}_d}{(\mu r)^d} \max_{z \in C_r} |f(z)| \| v \| \end{aligned}$$

Clearly, the above analysis assumes an idealized setting (e.g., in reality, the λ_i will not be equidistributed on ∂D) and additionally, the obtained bound cannot be expected to be sharp and descriptive of the actual difference of the two approximations. It should




FIGURE 1 | Example 4.3. True error $\| \exp(M_d + v_d e_d^T) e_1 - \exp(M_d) e_1 \|$ (dashed line) as the dimension d grows, and bound $(1 + \sqrt{2}) \| v_d \| \mathbf{m}_d / d!$ (solid line). Left: Toeplitz matrix. Right: M_d obtained by the Arnoldi recurrence with A from discretized PDE problem.

merely be seen as a first hint at what kind of results might be possible to obtain for general f . However, such more refined results will require advanced tools from complex analysis which are well beyond the scope of the current work.

5 | Error Analysis With Respect To the FOM Approximation

We can now use Theorem 4.1 to relate the different possible sketched and truncated Krylov approximations to each other, as they all result from rank-one modifications of one another. Note that for the sketched Krylov approximation, we will use the representation (17) for analysis purposes (as it directly uses the matrix H_d), although the representation (18) is better suited for actual computations, as we will illustrate in Example 7.1 below. Before stating our results, let us introduce yet another way of writing the sketched or truncated Krylov approximations: As both U_d and \mathcal{U}_d contain nested bases of $\mathcal{K}_d(A, b)$, there exists a nonsingular upper triangular matrix \mathcal{T}_d such that $U_d = \mathcal{U}_d \mathcal{T}_d$; the matrix \mathcal{T}_d can, for example, be obtained via a Gram–Schmidt-based QR decomposition. We write

$$\mathcal{T}_{d+1} = \begin{bmatrix} \mathcal{T}_d & t \\ 0^T & \hat{\tau}_{d+1} \end{bmatrix} \quad (30)$$

Using this, we can write the sketched FOM approximation in terms of the orthogonal Krylov basis \mathcal{U}_d as

$$f_d^{\text{SK}} = \mathcal{U}_d f(\mathcal{T}_d H_d \mathcal{T}_d^{-1} + \mathcal{T}_d r e_d^T \mathcal{T}_d^{-1}) \mathcal{T}_d e_1 \| b \| \quad (31)$$

Now, from (3) and the fact that the columns of \mathcal{U}_d are orthonormal it follows by straightforward algebraic manipulations that

$$\mathcal{T}_d H_d \mathcal{T}_d^{-1} = \mathcal{H}_d - h_{d+1,d} \mathcal{U}_d^T u_{d+1} e_d^T \mathcal{T}_d^{-1} \quad (32)$$

Thus, the matrix at which f is evaluated in (31) is a rank-one modification of \mathcal{H}_d . Further simplification is possible by noting that the first columns of U_d and \mathcal{U}_d are identical, so the top-left entry of \mathcal{T}_d can be chosen to be 1 and furthermore, $e_d^T \mathcal{T}_d^{-1} = e_d^T / \hat{\tau}_d$. Using these observations together with (32) allows to rewrite (31) as

$$f_d^{\text{SK}} = \mathcal{U}_d f(\mathcal{H}_d + \hat{\tau}_d e_d^T) e_1 \| b \| \quad (33)$$

where we introduced the shorthand notation

$$\begin{aligned} \hat{r} &= (\mathcal{T}_d r - h_{d+1,d} \mathcal{U}_d^T u_{d+1}) / \hat{\tau}_d \\ &= (\mathcal{T}_d T_d^{-1} t - \mathcal{U}_d^T u_{d+1}) (h_{d+1,d} / \hat{\tau}_d) \end{aligned} \quad (34)$$

$$= (\mathcal{T}_d T_d^{-1} t - t) (h_{d+1,d} / \hat{\tau}_d) \quad (35)$$

with t coming from (30).

By following the same steps for the truncated FOM approximation, we find the representation

$$f_d^{\text{TR}} = \mathcal{U}_d f(\mathcal{H}_d - h_{d+1,d} t e_d^T) e_1 \| b \| \quad (36)$$

Due to the many different possible formulations, we summarize all important matrix function approximations (both in the

truncated and fully orthonormal bases) in Table 1 to provide the reader with an overall picture. Among the various sketched forms we have discussed, we report here only those that appear to be most reliable for practical purposes.²

For studying the a priori error with respect to the ideal FOM approximation, writing the truncated and the sketched approximations in terms of the fully orthonormal basis \mathcal{U}_d (third column in Table 1) leads to more insightful formulas that we also exploit in the results below. It should be kept in mind, however, that in our setting the formulation of sketched and truncated approximations via the fully orthonormal basis is not computationally affordable.

Corollary 5.1. *The sketched and full Arnoldi approximations f_d^{SK} and f_d^{FOM} from (17) and (2) fulfill*

$$f_d^{\text{SK}} - f_d^{\text{FOM}} = \mathcal{U}_d g_{\hat{r}}(\mathcal{H}_d) \hat{r} \| b \| \quad (37)$$

where \hat{r} is defined in (35) and $g_{\hat{r}}$ is defined in (20), with respect to the eigendecomposition of $\mathcal{H}_d + \hat{\tau}_d e_d^T$.

Proof. The result follows by applying Theorem 4.1 with $M = \mathcal{H}_d$ and $v = \hat{r}$.

Corollary 5.2. *The truncated and full Arnoldi approximations f_d^{TR} and f_d^{FOM} from (4) and (2) fulfill*

$$f_d^{\text{TR}} - f_d^{\text{FOM}} = \mathcal{U}_d g_y(\mathcal{H}_d) y \| b \|$$

where $y = -h_{d+1,d} t$ and g_y is defined in (20), with respect to the eigendecomposition of $\mathcal{H}_d + y e_d^T$.

Proof. The result follows by applying Theorem 4.1 with $M = \mathcal{H}_d$ and $v = -h_{d+1,d} t$.

There are several possibilities for comparing the sketched and truncated Arnoldi approximations to each other. We present only the one based on the banded Hessenberg matrix H_d and the non-orthogonal basis U_d in the following. It would also be possible, using (33) and (36), to state a result in terms of the matrices \mathcal{U}_d and \mathcal{H}_d from the full Arnoldi process, but this is less appropriate here as neither of the two methods uses these matrices in practice.

Corollary 5.3. *The truncated and sketched Arnoldi approximations f_d^{TR} and f_d^{SK} from (4) and (17) fulfill*

$$f_d^{\text{TR}} - f_d^{\text{SK}} = U_d g_r(H_d) r \| b \|$$

where g_r is defined in (20), with respect to the eigendecomposition of $H_d + r e_d^T$.

Proof. The result follows from Theorem 4.1 applied with $M = H_d$ and $v = r$.

When bounding the distance between the different approximations, the results of Theorem 4.4 and Proposition 4.1 can be used to handle the matrix functions occurring on the right-hand sides of the preceding corollaries by bounding $\| g_v(\mathcal{H}_d) \|$. Such bounds on the norm of the matrix function are only useful if

we also have appropriate bounds available for the vectors v . In the following, we particularly discuss the vector $v = \hat{r}$ from Corollary 5.1 which relates the sketched to the full Arnoldi approximation.

Consider Equation (34). Using the fact that $t = Q_{d+1}^* S u_{d+1}$ together with $T_d^{-1} Q_d^* = (S U_d)^\dagger$, we find the representation

$$\begin{aligned}\hat{r} &= \left(\mathcal{F}_d (S U_d)^\dagger S u_{d+1} - \mathcal{U}_d^T u_{d+1} \right) (h_{d+1,d} / \hat{\tau}_d) \\ &= \mathcal{U}_d^T \left(U_d (S U_d)^\dagger S u_{d+1} - u_{d+1} \right) (h_{d+1,d} / \hat{\tau}_d)\end{aligned}$$

where we have used $\mathcal{F}_d = \mathcal{U}_d^T U_d$ for the second equality. The vector $(S U_d)^\dagger S u_{d+1}$ appearing above is the solution vector x_S of the sketched least squares problem

$$x_S = \arg \min_{x \in \mathbb{R}^d} \| S(U_d x - u_{d+1}) \|^2$$

From established theory covering sketching methods for least squares problems [2], it is known that the solution x_S of the sketched problem cannot have a significantly higher residual than the solution of the original problem if the embedding quality is high enough. Precisely, denoting by $x_\star = U_d^\dagger u_{d+1}$ the solution of the non-sketched least squares problem, we have the following inequality relating the residuals of x_S and x_\star ,

$$\| U_d x_S - u_{d+1} \| \leq \sqrt{\frac{1+\varepsilon}{1-\varepsilon}} \| U_d x_\star - u_{d+1} \|^2$$

see eq. (2.3) in [7], which is adapted from results in [2, 16].

Further, due to the fact that $U_d = \mathcal{U}_d \mathcal{F}_d$ is a QR decomposition, we have that $\| U_d x_\star - u_{d+1} \| = |\hat{\tau}_{d+1}|$. Since $\| U_d x_S - u_{d+1} \| = \|(I - \mathcal{U}_d \mathcal{U}_d^T) u_{d+1}\|$, the quantity $|\hat{\tau}_{d+1}|$ measures the contribution of the new vector u_{d+1} to the expansion of the Krylov subspace.

Putting all this together and using the fact that \mathcal{U}_d has orthonormal columns, we find

$$\| \hat{r} \| \leq h_{d+1,d} \left| \frac{\hat{\tau}_{d+1}}{\hat{\tau}_d} \right| \sqrt{\frac{1+\varepsilon}{1-\varepsilon}} \quad (38)$$

As long as the space keeps growing, we expect that $\left| \frac{\hat{\tau}_{d+1}}{\hat{\tau}_d} \right| \approx 1$, so that an overall result of the following form can be obtained from Proposition 4.1.

Corollary 5.4. *Let $\mathcal{H}_d \in \mathbb{C}^{d \times d}$ be the upper Hessenberg matrix from the (full) Arnoldi relation (1), $\mathbf{h}_d := \prod_{j=1}^{d-1} (\mathcal{H}_d)_{d+1,d}$, and let $D \subseteq \mathbb{C}$ be a convex set that contains $F(\mathcal{H}_d)$ and the eigenvalues of $\mathcal{H}_d + \hat{r} e_d^T$. Then*

$$\| f_d^{\text{SK}} - f_d^{\text{FOM}} \| \leq (1 + \sqrt{2}) \sqrt{\frac{1+\varepsilon}{1-\varepsilon}} \| b \| \frac{\mathbf{h}_d}{d!} \max_{z \in D} |f^{(d)}(z)| \quad (39)$$

The bound in Corollary 5.4 depends on the set D , and in general, D might be much larger and have less favorable properties than $F(\mathcal{H}_d)$ (which is guaranteed to lie within $F(A)$). For example, a straightforward choice of D would be the convex

hull of $F(\mathcal{H}_d) \cup F(\mathcal{H}_d + \hat{r} e_d^T)$. In that case, the bound (39) qualitatively reproduces available results from the literature, as, for example, theorem 4.3 in [5] and corollary 2.4 in [6]. These results rely on the Crouzeix-Palencia theorem applied directly to $f_d^{\text{SK}} - f_d^{\text{FOM}}$ and thus necessarily need to work with a region D in the complex plane that includes the field of values of the sketched and projected matrix. In contrast, we apply the Crouzeix-Palencia theorem only to h from (27), which is evaluated at \mathcal{H}_d . The influence of the rank-one modified matrix $\mathcal{H}_d + \hat{r} e_d^T$ thus only occurs through its eigenvalues, thus giving us more freedom in the choice of D . In particular, it allows us to use the results presented in Section 6 below to argue why eigenvalues far outside $F(\mathcal{H}_d)$ do not negatively influence the behavior of the method. This is in line with experimental evidence reported in [5, 6] as well as in our Example 7.1 below.

6 | On the Error Vector and the Role of Outlying Spurious Ritz Values

In the previous section, we wrote the error between different methods in terms of the “error” vector (see Theorem 4.1)

$$g_v(M)v = \sum_{i=1}^d \alpha_i \beta_i f[M, \lambda_i] v$$

where $\alpha_i = e_d^T X e_i$, $\beta_i = e_i^T X^{-1} e_1$ with X from the eigenvalue decomposition (19).

The magnitude of the α_i, β_i and of the divided differences needs further analysis. In general, and for large $\|v\|$, we do not expect the eigenvalues of \mathcal{H}_d and $\mathcal{H}_d + v e_d^T$ (or H_d and $H_d + v e_d^T$) to be close to each other. We now discuss this again in the particular situation of Corollary 5.1, that is, for $v = \hat{r}$, noting that the derivation works similarly in the other cases.

For any nonzero $d \times d$ lower Hessenberg matrix N , we define

$$\gamma_{\lambda, N} := \min_{i=1, \dots, d-1} \frac{|N_{i,i} - \lambda|}{|N_{i,i+1}|} \quad (40)$$

This quantity allows us to analyze the contribution of the eigenvalues of $\mathcal{H}_d + \hat{r} e_d^T$ in the error vector. We recall here that the eigenvalues of $\mathcal{H}_d + \hat{r} e_d^T$ are the same as those of $H_d + r e_d^T$. As soon as ill-conditioning arises in the generation of the basis U_d , the entries of r start to grow in magnitude. In this case, some of the eigenvalues λ_i may be far from the field of values of \mathcal{H}_d , as $F(\mathcal{H}_d + \hat{r} e_d^T)$ is not included in $F(\mathcal{H}_d)$, whereas others may be either close or within $F(\mathcal{H}_d)$. In the following, we show that whenever $\|r\|$ grows, the quantity β_i is positively affected by becoming correspondingly small.

We next assume that $\gamma_{\lambda_i, \mathcal{H}_d} > 1$. In particular, this is the case if the eigenvalue λ_i is far from $F(\mathcal{H}_d)$, relative to $\|\mathcal{H}_d\|$. It appears that a significantly distant eigenvalue does not interfere with the actual error, so that the corresponding term in the sum of g_v becomes negligible, in spite of possibly large values of $f[\mathcal{H}_d, \lambda]$. This argument is justified next.

Proposition 6.1. *Let N be a $d \times d$ lower Hessenberg matrix with $N_{k,k+1} \neq 0$, for all k , and let $\hat{N} = N + e_d v^*$ for some v . Then*

for each eigenpair (λ, y) of \hat{N} with $\|y\| = 1$ it holds, denoting the components $y = (\eta_1, \dots, \eta_d)^T$, that

$$|\eta_{k+1}| \leq \frac{|N_{k,k} - \lambda|}{|N_{k,k+1}|} |\eta_k| + \frac{\|N_{k,1:k-1}\|}{|N_{k,k+1}|}, \quad k = 1, \dots, d-1$$

and

$$|(N_{d,d} + v_d) - \lambda| \|\eta_d\| \leq \|v_{1:d-1} + N_{d,1:d-1}\|$$

Proof. Let (λ, y) be an eigenpair of \hat{N} , that is $\hat{N}y = \lambda y$. For each row k with $2 \leq k < d$ it holds

$$N_{k,k+1}\eta_{k+1} = -(N_{k,k} - \lambda)\eta_k + N_{k,1:k-1}\eta_{1:k-1}$$

We assume that not all indexed components of y are zero, otherwise the bound follows trivially. Hence,

$$\begin{aligned} |\eta_{k+1}| &\leq \frac{|N_{k,k} - \lambda|}{|N_{k,k+1}|} |\eta_k| + \frac{\|N_{k,1:k-1}\|}{|N_{k,k+1}|} \|\eta_{1:k-1}\| \\ &\leq \frac{|N_{k,k} - \lambda|}{|N_{k,k+1}|} |\eta_k| + \frac{\|N_{k,1:k-1}\|}{|N_{k,k+1}|} \end{aligned}$$

Equating the last row yields the second bound.

A few comments are in order. Let $\phi_k(\lambda) = \prod_{j=1}^k \frac{(N_{j,j} - \lambda)}{N_{j,j+1}}$ be the polynomial of degree k associated with the diagonal and super-diagonal of N .

If $\frac{|N_{j,j} - \lambda|}{|N_{j,j+1}|} > 1$ for each j , then the proposition above implies that

$$|\eta_{k+1}| \approx |\phi_k(\lambda)|, \quad k < d$$

This means that the components tend to grow as a power of k , that is

$$|\eta_k| \approx O(\gamma_{\lambda,N}^{k-1}) \quad (41)$$

Due to the unit norm of y , this means that if λ is sufficiently far from the diagonal elements of N , the first components of y are tiny, and the subsequent components grow as a growing power of such distance, with the last component of y being $O(1)$ (see Example 4.3). Finally, we observe that the inequality $\frac{|N_{j,j} - \lambda|}{|N_{j,j+1}|} > 1$ for each j may also hold if $|\lambda|$ is much larger than $|N_{j,j}|$ and $|N_{j,j+1}|$; in this case, the increasing property of the components is an intrinsic property of the eigenpair, and it may occur also for λ not far from the field of values of N .

The result of Proposition 6.1 is quite general. To apply it to our setting for the error formula, we take $N = \mathcal{H}_d^*$. Except for normalization, $y = y_i$ is the left eigenvector of \mathcal{H}_d corresponding to the eigenvalue λ_i for some i . In particular, $y_i^* e_1 = e_i^T X^{-1} e_1$. If λ_i is far away from the diagonal elements of \mathcal{H}_d^* , the contribution of $\beta_i = e_1^T X^{-1} e_i = (\bar{y}_i)_1$ becomes extremely limited, that is,

$$|\beta_i| \approx O(\gamma_{\lambda,N}^{1-d})$$

This property is crucial because it ensures the following counter-intuitive result: although the rank-one modification can dramatically affect the spectrum, the wildly varying eigenvalues are not going to contribute to the error.

The following example illustrates the eigenvector growth property.

Example 6.1. We consider the Toeplitz data in Example 4.3 for fixed $d = 100$, with \mathcal{H}_d^* playing the role of N and $(\mathcal{H}_d + \hat{r}e_d^T)^*$ that of $N + e_d v^T$. The left plot in Figure 2 shows the spectral quantities associated with the two matrices, including the boundary of $F(\mathcal{H}_d)$ for later reference. The plot clearly reports the presence of many eigenvalues of $\mathcal{H}_d + \hat{r}e_d^T$ outside the field of values of \mathcal{H}_d , and also falling into the right half-plane. The right plot shows the components absolute value of the two left eigenvectors of $\mathcal{H}_d + \hat{r}e_d^T$ corresponding to the eigenvalues pointed to in the left plot. The two additional lines report the values of $|\phi_k(\lambda)|$ for the two eigenvalues; the vectors containing these quantities

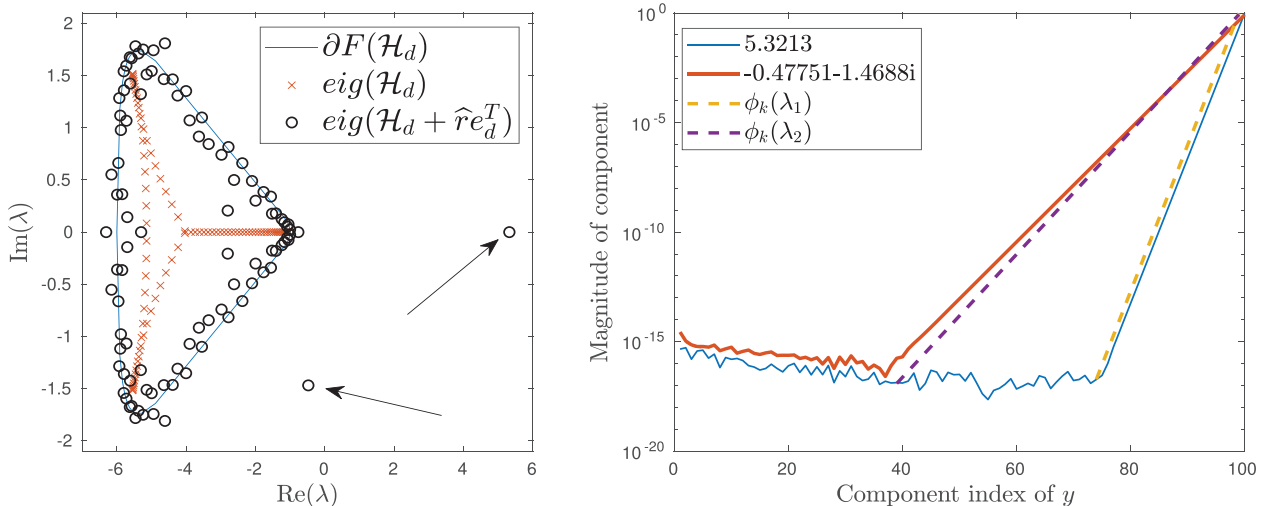


FIGURE 2 | Example 6.1. Left: Boundary of the field of values of \mathcal{H}_d (solid blue line) along with its eigenvalues of \mathcal{H}_d (red crosses) and the eigenvalues of $\mathcal{H}_d + \hat{r}e_d^T$ (black circles). The arrows point to the eigenvalues used in the right plot. Right: Absolute values of the eigenvector components for the selected eigenvalues, together with the normalized vector of polynomial values ϕ_k .

have been scaled so that the last component has unit value. The eigenvector components behavior is as expected, and it is accurately predicted by the quantity $|\phi_k(\lambda)|$. We notice that the first several components of each eigenvector have a magnitude at the level of machine epsilon. In exact arithmetic their values would follow the steep slope of the other components.

It is worth mentioning that whenever β_i is small, the term $\beta_i f[N, \lambda_i]$ remains small. For instance, for the exponential,

$$\beta_i \exp[N, \lambda_i] \approx \frac{|N_k|^{d-1}}{|N_{k,k+1} - \lambda_i|^{d-1}} (N - \lambda_i I)^{-1} (\exp(N) - \exp(\lambda_i) I)$$

for some $k < d$. Let us assume, as in the previous example, that $\gamma_{\lambda_i, N} \approx |\lambda_i| > 1$. Then

$$\frac{|\exp(\lambda_i)|}{(\gamma_{\lambda_i, N})^{d-1}} \approx \frac{|\exp(\lambda_i)|}{|\lambda_i|^{d-1}}$$

which goes to zero as d goes to infinity.

For our purposes, the analysis above shows that the likely erratic spectral behavior in the sketched approach will not negatively affect the final approximation. Loosely speaking, the matrix structure takes care of purging the undesired eigenvalues. In this way, the “bad” portion of the generated space is implicitly deflated.

Motivated by the analysis above, we return to the relation $\alpha_i \beta_i = \frac{\mathbf{m}_d}{\omega'(\lambda_i)}$. Let \mathcal{J} be the set of indexes i such that β_i is not tiny. Hence, the sum over all eigenvalues can be limited to those in \mathcal{J} , that is

$$\sum_{i=1}^d \alpha_i \beta_i f[M, \lambda_i] v \approx \sum_{i \in \mathcal{J}} \alpha_i \beta_i f[M, \lambda_i] v$$

A similar reduction can be written for (24), in terms of the d th order divided differences $f[M, \lambda_1, \dots, \lambda_d]$, namely $f[M, \lambda_i, \dots, \lambda_{i, \mathcal{J}_1}]$. Hence, the set \mathcal{J} identifies the “effective” eigenpairs on which the approximation error actually lives.

In the context of the sketched Arnoldi method, this shows that bounds like (39) would not be expected to be descriptive of the actual behavior of the method if they were to rely on approximation properties on the whole field of values $F(\mathcal{H}_d + \hat{r}e_d^T)$. Rather, the behavior of the method is dictated by a region that may be only slightly larger than the spectral region of \mathcal{H}_d . Precisely, in light of (26) and the analysis above, to obtain a convergence estimate for the sketched Arnoldi method, we can take D as any convex set containing $F(\mathcal{H}_d) \cup \{\lambda_i : i \in \mathcal{J}\}$. We therefore introduce a more descriptive region of the complex plane: Let $\delta \geq 0$ be such that all eigenvalues $\lambda_i, i \in \mathcal{J}$ satisfy $\text{dist}(\lambda, F(\mathcal{H}_d)) < \delta$, and let

$$F(\mathcal{H}_d, \delta) := \{z \in \mathbb{C} : \text{dist}(z, F(\mathcal{H}_d)) \leq \delta\}$$

we will call this the *effective* field of values of \mathcal{H}_d . Clearly, $F(\mathcal{H}_d) \subseteq F(\mathcal{H}_d, \delta)$.

For λ_i with $i \notin \mathcal{J}$, we let $\zeta_i = \arg \min_{z \in \partial F(\mathcal{H}_d)} |z - \lambda_i|$. Then we define the following set of values

$$\hat{\lambda}_i = \begin{cases} \lambda_i & i \in \mathcal{J} \\ \zeta_i & i \notin \mathcal{J} \end{cases}$$

Then, according to the previous discussion,

$$\sum_{i=1}^d \frac{1}{\omega'(\lambda_i)} f[\mathcal{H}_d, \lambda_i] v \approx \sum_{i=1}^d \frac{1}{\omega'(\hat{\lambda}_i)} f[\mathcal{H}_d, \hat{\lambda}_i] v$$

The definition of the $\hat{\lambda}_i$'s thus allows us to move (rather than remove) the inactive eigenvalues, while maintaining the number of terms in the divided difference. With these definitions, and using $h(z) := f[z, \hat{\lambda}_1, \dots, \hat{\lambda}_d]$, (29), applied for $M = \mathcal{H}_d$, becomes

$$\begin{aligned} \|h(\mathcal{H}_d)\| &\leq (1 + \sqrt{2}) \max_{z \in F(\mathcal{H}_d)} |h(z)| \\ &\leq (1 + \sqrt{2}) \frac{\max_{\zeta \in F(\mathcal{H}_d, \delta)} |f^{(d)}(\zeta)|}{d!} \end{aligned}$$

The bound employs the region of the complex plane where the sketched method is active, and this region can be much smaller than the field of values of $\mathcal{H}_d + \hat{r}e_d^T$. By means of these new tools, we can obtain an approximate bound of the form

$$\begin{aligned} &\|f(\mathcal{H}_d + \hat{r}e_d^T)e_1 - f(\mathcal{H}_d)e_1\| \\ &\leq \mathbf{m}_d (1 + \sqrt{2}) \frac{\max_{\zeta \in F(\mathcal{H}_d, \delta)} |f^{(d)}(\zeta)|}{d!} \|\hat{r}\| \end{aligned} \quad (42)$$

which is a family of analogues of (26) depending on the effective field of values, or more precisely on δ .

To summarize, our analysis shows that convergence of the sketched method will always depend on a region that is not much larger than that of \mathcal{H}_d (and thus that of A): If sketched Ritz values far outside the spectral region of \mathcal{H}_d occur, they simply do not play a role. Only those that are slightly outside need to be accounted for. Thus, one can typically expect at most a slight delay of convergence due to sketching. The role of $\gamma_{\lambda, N}$ is crucial, and its relation with the field of values of N deserves further work, which will be postponed to a later investigation: while $\gamma_{\lambda, N} > 1$ whenever λ is far away from $F(\mathcal{H}_d)$, we also observed that this condition might sometimes hold for eigenvalues close to or even within the field of values, indicating that there might be even more sketched Ritz values that do not significantly influence the error.

A by-product of this analysis is the following: If none of the sketched Ritz values lie outside of $F(\mathcal{H}_d)$, then we can take $\delta = 0$ above, and the convergence behavior of the sketched Arnoldi method simply depends on properties of f on $F(\mathcal{H}_d, 0) = F(\mathcal{H}_d)$, irrespective of how large $F(\hat{H}_d + \hat{r}e_d^T)$ is. This is in contrast to, for example, the results of [5] which rely on the Crouzeix-Palencia theorem [25] and therefore always depend on the whole field of values of $\hat{H}_d + \hat{r}e_d^T$, even if the sketched Ritz values are well-behaved.

7 | Robust Computation of Sketched Approximations

In this section, we want to briefly discuss two further issues related to the experimentally observed numerical (in)stability of sketched Krylov methods. On the one hand, we provide arguments that explain why the WS-Arnoldi version (18) of the sFOM approximation typically appears to be very robust in floating point arithmetic despite the inversion of the possibly very ill-conditioned matrix T_d that is involved. On the other hand, we explain which potential problems can occur with the non-whitened variant (17) based on a rank-one modification of H_d .

By starting from expression (18) for the sketched-whitened approximation, that is $f_d^{\text{SK}} = U_d T_d^{-1} \hat{y}$ with $\hat{y} = f(\hat{H}_d + \hat{t}e_d^T)e_1 \|Sb\|$, we want to analyze the stability properties of $T_d^{-1} \hat{y}$ and show that the solution accuracy is not significantly influenced by the possible ill-conditioning of T_d . To this end, we recall the well-known fact that in general, triangular systems are often solved to much higher accuracy in floating point arithmetic than their condition number would suggest (although it is not possible to prove a precise, general results in this direction as—rather academic—counter examples exist). This observation goes back to Wilkinson [30] and has since then been investigated many times; see, for example, Chapter 8 in [31] for an overview. In the following, we provide a result that is specifically tailored to our situation.

We first recall the following standard forward error bound for triangular systems; see, for example, Equation (2.6) in [32]. Let $z = T_d^{-1} \hat{y}$ and \tilde{z} be the exact and computed solutions, respectively. Then

$$\frac{\|z - \tilde{z}\|_\infty}{\|z\|_\infty} \leq (d + 1) u \text{cond}(T_d, z) + O(u^2) \quad (43)$$

where

$$\text{cond}(T_d, z) = \frac{\| |T_d^{-1}| |T_d| |z| \|_\infty}{\|z\|_\infty}$$

and $|\cdot|$ is the component-wise absolute value. Here u is the unit roundoff.

Proposition 7.1. *Let $\text{cond}_\infty(T) = \|T\|_\infty \|T^{-1}\|_\infty$ denote the infinity norm condition number. Assume T_d can be partitioned in such a way that*

$$T_d = \begin{bmatrix} T_1 & T_2 \\ 0 & \tau T_3 \end{bmatrix}, \quad \text{cond}_\infty(T_1) = c_1, \quad \text{and} \quad \text{cond}_\infty(T_3) = c_3$$

and $\tau > 0$. Both T_1, T_3 are square and nonsingular. Then (43) holds with

$$\text{cond}(T_d, z) \leq \frac{\max\{c_3 \|z_2\|_\infty, c_1 \|z_1\|_\infty + \| |T_1^{-1} T_2| \|_\infty (1 + c_3) \|z_2\|_\infty\}}{\|z\|_\infty}$$

where $z = [z_1; z_2]$ is partitioned accordingly with T_d .

Proof. Let $\tilde{T}_3 = \tau T_3$. We spell out the definition of $\text{cond}(T_d, z)$ in terms of the partitioning of T_d . We first write $T_d^{-1} = \begin{bmatrix} T_1^{-1} & -T_1^{-1} T_2 \tilde{T}_3^{-1} \\ 0 & \tilde{T}_3^{-1} \end{bmatrix}$. We then have

$$\begin{aligned} \| |T_d^{-1}| |T_d| |z| \|_\infty &= \| |T_d^{-1}| \begin{bmatrix} |T_1| |z_1| + |T_2| |z_2| \\ |\tilde{T}_3| |z_2| \end{bmatrix} \|_\infty \\ &= \left\| \begin{bmatrix} |T_1^{-1}| (|T_1| |z_1| + |T_2| |z_2|) + |T_1^{-1} T_2 \tilde{T}_3^{-1}| |\tilde{T}_3| |z_2| \\ | \tilde{T}_3^{-1} | |\tilde{T}_3| |z_2| \end{bmatrix} \right\|_\infty \\ &\leq \left\| \begin{bmatrix} c_1 \|z_1\|_\infty + \| |T_1^{-1} T_2| \|_\infty (1 + c_3) \|z_2\|_\infty \\ c_3 \|z_2\|_\infty \end{bmatrix} \right\|_\infty \\ &= \max\{c_3 \|z_2\|_\infty, c_1 \|z_1\|_\infty + \| |T_1^{-1}| |T_2| \|_\infty (1 + c_3) \|z_2\|_\infty\} \end{aligned}$$

□

Proposition 7.1 shows that, as long as the ill-conditioning of T_d is encoded in its right-bottom corner as the multiplicative term τ , the entry-wise condition number $\text{cond}(T_d, z)$, and thus the forward error in (43), is bounded solely by the condition numbers of well-behaved submatrices and exact arithmetic quantities. Moreover, notice that T_3 in Proposition 7.1 does not need to be a matrix but it can also be just a scalar. In this case it holds $c_3 = 1$ and $\tau = \tau_d$.

The assumptions in Proposition 7.1 are easily met in our sketched-and-truncated framework. Indeed, the non-orthogonal basis U_d gradually becomes more ill-conditioned as the iteration progresses and (almost) linear dependence occurs mostly between “later” basis vectors. To understand this, consider the extreme case of no orthogonalization at all: Then, the sequence of Krylov basis vectors converges towards the dominant eigenvector of A , so that later basis vectors will all point in almost the same direction, while staying well linearly independent with respect to the first basis vectors. A similar phenomenon is preserved when employing a partial orthogonalization with small truncation parameter.

As the sketching matrix S distorts inner products in a controlled manner due to (6), this property of U_d is inherited by the sketched basis SU_d for which we perform the QR decomposition. This observation implies that the severe ill-conditioning of T_d is mostly encoded in a bottom right subblock, the same exact scenario depicted in Proposition 7.1.

Remark 3. Considering (17), note that the matrix $H_d + re_d^T$ is often not well-balanced, in the sense that the norm of its last column is typically much larger than that of all other columns. It is known since the 1970s that balancing can have an influence on the approximation accuracy for matrix functions; for example, when using techniques such as Padé approximation [33, 34], although the effect is difficult to rigorously quantify; cf., for example, sect. 3 and 6 in [35].

We turn to the specific case of scaling-and-squaring methods for the matrix exponential. A first step in such methods is the computation of a parameter γ such that $\|2^{-\gamma} (H_d + re_d^T)\| < 1$. When $H_d + re_d^T$ is poorly balanced, the value of γ will be almost solely determined by its last column, and the entries of all other columns in $2^{-\gamma} (H_d + re_d^T)$ might become tiny. This and related

effects are known under the term *overscaling* in the matrix function literature; see, for example, sect. 1 in [36] or [37]. Clearly, this can result in a noticeable decrease in the accuracy of the method in the presence of round-off error. In contrast, the matrix $\hat{H}_d + \hat{t}e_d^T$ from (17) will typically be better balanced due to the similarity transformation with T_d and thus less prone to overscaling effects.

We next report the results of an illustrative example. For the sake of clarity we report an algorithmic description of the whitened-sketched Arnoldi method that we use as Algorithm 1 in Appendix A. We stress that in our implementation the QR decomposition employed in the whitening procedure was updated from one iteration to the next instead of recomputing it, therefore causing only a modest computational cost increase, for s significantly smaller than n .

Example 7.1. In this example, A is obtained by discretizing the convection-diffusion operator

$$\mathcal{L}(u) = -\nu\Delta u + \vec{w} \cdot u \tag{44}$$

on the unit square $[0,1]^2$ by centered finite differences. We use viscosity parameter $\nu = 10^{-2}$, the convection field $\vec{w} = \left(\frac{3}{2}y(1-x^2), -3x(1-y^2)\right)$, and $N = 50$ discretization points in each spatial direction. The resulting matrix $A \in \mathbb{R}^{N^2 \times N^2}$ is highly non-normal with extremely ill-conditioned eigenvector basis, $\kappa(X) \approx 1.5 \cdot 10^{22}$. We aim to approximate the action of the matrix exponential $\exp(-A)b$, where b is the normalized vector of all ones.

We begin by comparing different possible implementations of sFOM, using a truncation parameter $k = 2$ in the generation of the non-orthogonal Krylov basis U_d and a sketching dimension $s = 400$. Specifically, we compare the whitened-sketched sFOM approximation (18) with the version (17) based on the rank one

modification $H_d + re_d^T$. Due to the bad balancing of $H_d + re_d^T$, the computation of $\exp(H_d + re_d^T)$ might be extremely sensitive to round-off, also depending on which algorithm is used for the matrix exponential; cf. also the discussion in Remark 3. We therefore compare two different versions: one which directly uses the built-in MATLAB function `expm` (which implements the scaling-and-squaring method from [36]) and one which first computes a Schur decomposition of A and then applies `expm` only to the triangular factor. Note that this can lead to better numerical stability for two reasons: First, the triangular factor of the Schur decomposition will in general be better balanced than $H_d + re_d^T$, and second, `expm` contains several implementation tricks to specifically improve the behavior for triangular inputs; see sect. 2 in [36]. For the whitened-sketched Arnoldi method, due to the better balancing, using either `expm` directly or working with a Schur decomposition gives virtually identical results so that we do not report separate curves here.

The results of this experiment are depicted on the left-hand side of Figure 3. It can be clearly observed that the whitened-sketched method closely follows the convergence of the full Arnoldi method, with only a very small deviation towards the end of the process. In contrast, both versions of sFOM based on the rank-one modification (17) show signs of instability. While all sketched methods behave identical for the first 90 iterations, the version based on `expm` becomes very unstable after that and shows increasing/stagnating error norms for about 50 iterations before continuing to converge after that. The version based on the Schur decomposition longer behaves similarly to the whitened-sketched method, but also starts to deviate around the 130th iteration, when reaching a relative error norm of about 10^{-8} .

It is a common conception (see, e.g., [7],) that convergence of sketched methods typically deteriorates once the non-orthogonal




FIGURE 3 | *Left:* Comparison of different versions of the sFOM approximation with full Arnoldi for approximating the action of the matrix exponential. The conditioning of the non-orthogonal Krylov basis U_d is also depicted. *Right:* Comparison of “best” sFOM implementation (i.e., whitened-sketched) with truncated Arnoldi. In both cases, the truncation parameter is $k = 2$ and the sketching parameter is $s = 400$.

basis becomes numerically linearly dependent (i.e., once it has a condition number larger than $\varepsilon_{\text{machine}}^{-1}$, the reciprocal of the unit round-off). Our experiment clearly reveals that this is not necessarily true and that there are also other factors which have an influence. For sFOM using expm, stability problems occur long before the basis becomes severely ill-conditioned, and the whitened-sketched method continues to converge smoothly also once this happens. Only for sFOM using a Schur decomposition, the linear dependence of the basis and occurrence of stability problems coincide.

In the right part of Figure 3, we compare the best performing of the sFOM variants, that is, the whitened-sketched version (18), to truncated Arnoldi (4) using the same truncation parameter $k = 2$. We can observe that in truncated Arnoldi, convergence is significantly delayed. In the initial 110 iterations of the method, a plateau of stagnating (or even slightly increasing) error norms occurs. After that, convergence begins, and the convergence rate is only slightly slower than that of the full (and whitened-sketched) method. The truncated method reaches the target relative error norm 10^{-11} after 200 iterations. Thus, the sketched-whitened method uses 25% fewer matrix–vector products in this case.

8 | Conclusions

We have provided a new analysis of sketched Krylov subspace methods. In particular, we have derived a new *sketched Arnoldi relation*, which applies to any Krylov method that combines subspace embeddings with a truncated orthogonalization process. This relation can be used to obtain new insights into the behavior of these methods and better explains some phenomena that are frequently observed, for example, that methods continue to perform well even if they produce spurious Ritz values far outside the field of values of A .

To demonstrate this, we have specifically focused on the application of sketching to the approximation of $f(A)b$, the action of a matrix function on a vector. We have derived new formulas for expressing certain rank-one modifications of a matrix function via divided differences and then used those to compare sketched and truncated Krylov approximations to the standard (full) Arnoldi approximation. In particular, we could prove that for functions like the exponential, for growing Krylov dimension d , the sketched Arnoldi approximation is guaranteed to converge to the full Arnoldi approximation, and the nature of our bound suggests that one can expect this convergence to take place at roughly the same speed as convergence of the full Arnoldi approximation to $f(A)b$.

Our focus in this work was on deepening the theoretical understanding of existing sketching approaches, not on introducing new algorithmic concepts. Therefore, we did not perform extensive, large-scale numerical experiments, as thorough experimental studies illustrating the potential of sketched Krylov methods for matrix functions already exist in the literature; see, for example [5, 6].

In [38] we employ the sketched Arnoldi approximation derived here for analyzing and implementing sketched and truncated Krylov subspace methods for efficiently solving matrix Sylvester and Lyapunov equations.

Acknowledgments

The first and third authors are members of the INdAM Research Group GNCS that partially supported this work through the funded project GNCS2022 “Tecniche Avanzate per Problemi Evolutivi: Discretizzazione, Algebra Lineare Numerica, Ottimizzazione” (CUP_E55F22000270001).

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

The data that support the findings of this study are available from the corresponding author upon reasonable request.

Endnotes

¹ Forming this matrix would incur a cost of $\mathcal{O}(d^2n)$ and thus be as expensive as performing the full Arnoldi method in the first place, so this needs to be avoided.

² For a matrix analysis, the truncated-Arnoldi form of the sketched approximation, namely $f_d^{\text{SK}} = U_d f(H_d + re_d^T) e_1 \|b\|$, would have several advantages. However, as we illustrate in Example 7.1 this form may showcase a quite erratic numerical behavior, compared with the stabilized one reported in the table.

References

1. N. Halko, P. G. Martinsson, and J. A. Tropp, “Finding Structure With Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions,” *SIAM Review* 53, no. 2 (2011): 217–288.
2. T. Sarlós, “Improved Approximation Algorithms for Large Matrices via Random Projections,” in *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS’06)* (Berkeley, CA: IEEE, 2006), 143–152.
3. F. Woolfe, E. Liberty, V. Rokhlin, and M. Tygert, “A Fast Randomized Algorithm for the Approximation of Matrices,” *Applied and Computational Harmonic Analysis* 25 (2008): 335–366.
4. O. Balabanov and L. Grigori, “Randomized Gram–Schmidt Process With Application to GMRES,” *SIAM Journal on Scientific Computing* 44, no. 3 (2022): A1450–A1474.
5. 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, no. 1 (2024): 619–633.
6. S. Güttel and M. Schweitzer, “Randomized Sketching for Krylov Approximations of Large-Scale Matrix Functions,” *SIAM Journal on Matrix Analysis and Applications* 44, no. 3 (2023): 1073–1095.
7. 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, no. 2 (2024): 1183–1214.
8. C. Lanczos, “An Iteration Method for the Solution of the Eigenvalue Problem of Linear Differential and Integral Operators,” *Journal of Research of the National Institute of Standards and Technology* 45 (1950): 255–282.
9. O. Balabanov and A. Nouy, “Randomized Linear Algebra for Model Reduction. Part I: Galerkin Methods and Error Estimation,” *Advances in Computational Mathematics* 45, no. 5 (2019): 2969–3019.

10. E. Timsit, L. Grigori, and O. Balabanov, “Randomized Orthogonal Projection Methods for Krylov Subspace Solvers,” *arXiv preprint arXiv:2302.07466* (2023).
11. The MathWorks, Inc, “*MATLAB 7*,” r2020b ed (2020).
12. W. E. Arnoldi, “The Principle of Minimized Iteration in the Solution of the Matrix Eigenvalue Problem,” *Quarterly of Applied Mathematics* 9 (1951): 17–29.
13. E. Carson, K. Lund, M. Rozložník, and S. Thomas, “Block Gram–Schmidt Algorithms and Their Stability Properties,” *Linear Algebra and its Applications* 638 (2022): 150–195.
14. Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd ed. (Philadelphia: SIAM, 2000).
15. P. Drineas, M. W. Mahoney, and S. Muthukrishnan, *Subspace Sampling and Relative-Error Matrix Approximation: Column-Based Methods* (Berlin Heidelberg: Springer, 2006), 316–326.
16. D. P. Woodruff, “Sketching as a Tool for Numerical Linear Algebra,” *Foundations and Trends in Theoretical Computer Science* 10, no. 1–2 (2014): 1–157.
17. P. G. Martinsson and J. A. Tropp, “Randomized Numerical Linear Algebra: Foundations and Algorithms,” *Acta Numerica* 29 (2020): 403–572.
18. J. A. Tropp, “Improved Analysis of the Subsampled Randomized Hadamard Transform,” *Advances in Adaptive Data Analysis* 3, no. 01n02 (2011): 115–126.
19. C. Sohler and D. P. Woodruff, “Subspace Embeddings for the L1-Norm With Applications,” in *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing* (San Jose, CA: IEEE, 2011), 755–764.
20. B. Philippe and L. Reichel, “On the Generation of Krylov Subspace Bases,” *Applied Numerical Mathematics* 62, no. 9 (2012): 1171–1186.
21. M. Ilić, I. W. Turner, and D. P. Simpson, “A Restarted Lanczos Approximation to Functions of a Symmetric Matrix,” *IMA Journal of Numerical Analysis* 30, no. 4 (2010): 1044–1061.
22. C. de Boor, “Divided Differences,” *Surveys in Approximation Theory* 1 (2005): 46–69.
23. M. Eiermann and O. G. Ernst, “A Restarted Krylov Subspace Method for the Evaluation of Matrix Functions,” *SIAM Journal on Numerical Analysis* 44, no. 6 (2006): 2481–2504.
24. H. Tal-Ezer, “On Restart and Error Estimation for Krylov Approximation of $w = f(A)v$,” *SIAM Journal on Scientific Computing* 29, no. 6 (2007): 2426–2441.
25. M. Crouzeix and C. Palencia, “The Numerical Range Is a $(1 + \sqrt{2})$ -Spectral Set,” *SIAM Journal on Matrix Analysis and Applications* 38, no. 2 (2017): 649–655.
26. N. J. Higham, *Functions of Matrices: Theory and Computation* (Philadelphia: SIAM, 2008).
27. R. A. Horn and C. R. Johnson, *Topics in Matrix Analysis* (Cambridge: Cambridge University Press, 1991).
28. E. Gallopoulos and Y. Saad, “Efficient Solution of Parabolic Equations by Krylov Approximation Methods,” *SIAM Journal on Scientific and Statistical Computing* 13, no. 5 (1992): 1236–1264.
29. J. H. Curtiss, “Limits and Bounds for Divided Differences on a Jordan Curve in the Complex Domain,” *Pacific Journal of Mathematics* 12, no. 4 (1962): 1217–1233.
30. J. H. Wilkinson, “Error Analysis of Direct Methods of Matrix Inversion,” *Journal of the ACM* 8, no. 3 (1961): 281–330.
31. N. J. Higham, *Accuracy and Stability of Numerical Algorithms* (Philadelphia, PA: SIAM, 2002).
32. N. J. Higham, “Stability of Parallel Triangular System Solvers,” *SIAM Journal on Scientific Computing* 16, no. 2 (1995): 400–413.
33. W. Fair and Y. L. Luke, “Padé Approximations to the Operator Exponential,” *Numerische Mathematik* 14, no. 4 (1970): 379–382.
34. R. C. Ward, “Numerical Computation of the Matrix Exponential With Accuracy Estimate,” *SIAM Journal on Numerical Analysis* 14, no. 4 (1977): 600–610.
35. A. H. Al-Mohy and N. J. Higham, “Computing the Action of the Matrix Exponential, With an Application to Exponential Integrators,” *SIAM Journal on Scientific Computing* 33, no. 2 (2011): 488–511.
36. A. H. Al-Mohy and N. J. Higham, “A New Scaling and Squaring Algorithm for the Matrix Exponential,” *SIAM Journal on Matrix Analysis and Applications* 31, no. 3 (2010): 970–989.
37. L. Dieci and A. Papini, “Padé Approximation for the Exponential of a Block Triangular Matrix,” *Linear Algebra and its Applications* 308, no. 1–3 (2000): 183–202.
38. D. Palitta, M. Schweitzer, and V. Simoncini, “Sketched and Truncated Polynomial Krylov Methods: Matrix Sylvester Equations,” (Math. Comp., 2024).
39. B. Beckermann, D. Kressner, and M. Schweitzer, “Low-Rank Updates of Matrix Functions,” *SIAM Journal on Matrix Analysis and Applications* 39, no. 1 (2018): 539–565.

Appendix A

Pseudocode for Sketched-and-Truncated Arnoldi Method

We report here the pseudocode illustrating the sketch-and-truncated Krylov subspace method for the approximation of $f(A)b$.

We stress that adopting the whitening (lines 8 and 9 of Algorithm 1) does not remarkably increase the overall computational cost of the procedure. Indeed, assuming the sketched basis SU_d and its QR factors Q_d and T_d have been stored (these are small dimensional matrices), line 8 requires the application of the sketching to u_{d+1} , whose cost depends on the selected sketching but it is often only polylogarithmic in n , and the update of the QR factorization of $[SU_d, Su_{d+1}]$, which can be computed in $O(sd)$ flops by, for example, a Gram-Schmidt procedure. Moreover, we can fully take advantage of the upper Hessenberg, banded structure of H_d and the triangular pattern of T_d to reduce the cost of updating \hat{H}_d in line 9; see sect. 5 in [38].

Notice that Algorithm 1 can be additionally enhanced with a two-pass strategy to avoid storing the whole basis U_d . If this strategy is implemented, step 10 is postponed to after the end of the loop, and step 11 is reduced to store only the last $k+1$ vectors. Finally, an ad-hoc stopping criterion tailored to the function f at hand can be included in the for-loop.

ALGORITHM 1 | Sketched-and-truncated Krylov method for $f(A)b$.

Input: $A \in \mathbb{R}^{n \times n}$, $b \in \mathbb{R}^n$, $S \in \mathbb{R}^{s \times n}$, integers $0 < k < \maxit \ll n$

Output: $f_d^{\text{SK}} \approx f(A)b$

1: Set $U_1 = u_1 = b / \|b\|$, $Q_1 = SU_1 / \|SU_1\|$, $T_1 = \|SU_1\|$

2: **for** $d = 1, \dots, \maxit$ **do**

3: Compute $\tilde{u} = Au_d$

4: **for** $i = \max\{1, d - k + 1\}, \dots, d$ **do**

5: Set $\tilde{u} = \tilde{u} - u_i h_{i,d}$, where $h_{i,d} = \tilde{u}^T u_i$

6: **end for**

7: Set $h_{d+1,d} = \|\tilde{u}\|$ and $u_{d+1} = \tilde{u} / h_{d+1,d}$

8: Update skinny QR: $Q_{d+1}T_{d+1} = [SU_d, Su_{d+1}]$

9: Update $\hat{H}_d = T_d H_d T_d^{-1}$ and set $\hat{t} = (h_{d+1,d} / \tau_d) t$

10: Compute $f_d^{\text{SK}} = U_d (T_d^{-1} (f(\hat{H}_d + \hat{t} e_d^T) e_1) \|Sb\|)$

11: Set $U_{d+1} = [U_d, u_{d+1}]$

12: **end for**

Appendix B

Proofs

In this section, we report the proofs of Theorem 4.1 and Theorem 4.2.

Proof of Theorem 4.1. We follow the steps in Theorem 2.1 of [21]. We first write

$$\begin{aligned} & (zI - M - ve_d^T)^{-1}e_1 - (zI - M)^{-1}e_1 \\ &= (zI - M)^{-1} \left[(zI - M)(zI - M - ve_d^T)^{-1}e_1 - e_1 \right] \\ &= (zI - M)^{-1}v \left[e_d^T(zI - M - ve_d^T)^{-1}e_1 \right] \end{aligned}$$

Using the eigendecomposition of $M + ve_d^T$, we have

$$e_d^T(zI - M - ve_d^T)^{-1}e_1 = \sum_{i=1}^d \frac{\alpha_i \beta_i}{z - \lambda_i}$$

Let Γ be a closed curve enclosing the eigenvalues of $M + ve_d^T$ and of M . Then

$$\begin{aligned} & f(M + ve_d^T)e_1 - f(M)e_1 \\ &= \frac{1}{2\pi i} \int_{\Gamma} f(z) \left((zI - M - ve_d^T)^{-1}e_1 - (zI - M)^{-1}e_1 \right) dz \quad (B1) \\ &= \frac{1}{2\pi i} \int_{\Gamma} f(z)(zI - M)^{-1}v \left(e_d^T(zI - M - ve_d^T)^{-1}e_1 \right) dz \end{aligned}$$

$$\begin{aligned} &= \frac{1}{2\pi i} \sum_{i=1}^d \alpha_i \beta_i \int_{\Gamma} f(z) \frac{1}{z - \lambda_i} (zI - M)^{-1} dz \ v \\ &= \sum_{i=1}^d \alpha_i \beta_i f[M, \lambda_i] v \quad (B2) \\ &= g_v(M)v \end{aligned}$$

where we used the contour integral representation of divided differences (see, e.g., eq. 51 in [22]) in the second to last equality.

Remark 4. We note that in lemma 2.2 of [39] a conceptually different representation for rank-one updates of matrix functions was found—also based on the integral representation (B1)—but that this representation is less useful for our analysis.

We then proceed with showing Theorem 4.2.

Proof of Theorem 4.2. Following the proof of Theorem 4.1, formula (B1) dwells with the quantity

$$\mathcal{E} := \int_{\Gamma} f(z)(zI - M)^{-1}v \left(e_d^T(zI - M - ve_d^T)^{-1}e_1 \right) dz \quad (B3)$$

Now, let $\hat{M} = M + ve_d^T$ and $\hat{M}_z = zI - \hat{M}$, where $z \in \mathbb{C}$ is such that $zI - \hat{M}$ is nonsingular. Further denote by $\phi_d(z) = z^d + c_{d-1}z^{d-1} + \dots + c_1z + c_0$ the characteristic polynomial of \hat{M}_z , so that by the Cayley-Hamilton theorem, we have $\phi_d(\hat{M}_z) = \hat{M}_z^d + c_{d-1}\hat{M}_z^{d-1} + \dots + c_1\hat{M}_z + c_0I = 0$, with $c_0 = (-1)^d \det(\hat{M}_z)$. Then we have

$$\hat{M}_z^{-1} = \frac{-1}{c_0} (\hat{M}_z^{d-1} + c_{d-1}\hat{M}_z^{d-2} + \dots + c_1I_d)$$

Since M , and thus \hat{M}_z , is upper Hessenberg, it holds

$$\begin{aligned} e_d^T \hat{M}_z^{-1} e_1 &= \frac{(-1)^{d+1}}{\det(\hat{M}_z)} e_d^T \hat{M}_z^{d-1} e_1 = \frac{(-1)^{d+1}}{\det(\hat{M}_z)} \prod_{j=1}^{d-1} \hat{M}_{j+1,j} \\ &= \frac{(-1)^{2d}}{\det(\hat{M}_z)} \prod_{j=1}^{d-1} M_{j+1,j} = \frac{1}{\det(\hat{M}_z)} \prod_{j=1}^{d-1} M_{j+1,j} \end{aligned}$$

Moreover,

$$\det(\hat{M}_z) = \prod_{j=1}^d \lambda_j(\hat{M}_z) = \prod_{j=1}^d (z - \lambda_j)$$

where λ_j are the eigenvalues of \hat{M} . Therefore,

$$e_d^T(zI - M - ve_d^T)^{-1}e_1 = \prod_{j=1}^{d-1} M_{j+1,j} \prod_{i=1}^d (z - \lambda_i)^{-1}$$

We can thus write the integral in (B3) as

$$\mathcal{E} = \prod_{j=1}^{d-1} M_{j+1,j} \int_{\Gamma} \frac{f(z)}{\omega(z)} (zI - M)^{-1}v dz$$

To obtain (23), we write the partial fraction expansion of $1/\omega(z)$, that for distinct eigenvalues is given by

$$\frac{1}{\omega(z)} = \sum_{i=1}^d \frac{1}{\omega'(\lambda_i)} \frac{1}{z - \lambda_i}$$

where the prime denotes differentiation of $\omega(z)$ with respect to z , yielding

$$\int_{\Gamma} \frac{f(z)}{\omega(z)} (zI - M)^{-1}v dz = \sum_{i=1}^d \frac{1}{\omega'(\lambda_i)} \int_{\Gamma} \frac{f(z)}{z - \lambda_i} (zI - M)^{-1}v dz$$

Using once again the contour integral representation of divided differences and collecting the other terms, the formula (23) is proved.

To obtain (24) we use the integral representation (22) to write

$$\int_{\Gamma} \frac{f(z)}{\omega(z)} (zI - M)^{-1} dz = f[M, \lambda_1, \dots, \lambda_d]$$

or, equivalently, we observe that

$$\frac{1}{2\pi i} \sum_{i=1}^d \frac{1}{\omega'(\lambda_i)} f[z, \lambda_i] = f[z, \lambda_1, \dots, \lambda_d]$$

see, for example, eq. (1.1) in [29].