Tensor-Krylov method for computing eigenvalues of parameter-dependent matrices

Koen Ruymbeek^{a,*}, Karl Meerbergen^a, Wim Michiels^a

^a*Department of Computer Science, KU Leuven, Celestijnenlaan 200A, 3001 Leuven, Belgium*

Abstract

In this paper we extend the Residual Arnoldi method for calculating an extreme eigenvalue (e.g. largest real part, dominant,...) to the case where the matrices depend on parameters. The difference between this Arnoldi method and the classical Arnoldi algorithm is that in the former the residual is added to the subspace. We develop a Tensor-Krylov method that applies the Residual Arnoldi method (RA) for a grid of parameter points at the same time. The subspace contains an approximate Krylov space for all these points. Instead of adding the residuals for all parameter values to the subspace we create a low-rank approximation of the matrix consisting of these residuals and add only the column space to the subspace. In order to keep the computations efficient, it is needed to limit the dimension of the subspace and to restart once the subspace has reached the prescribed maximal dimension. The novelty of this approach is twofold. Firstly, we observed that a large error in the low-rank approximations is allowed without slowing down the convergence, which implies that we can do more iterations before restarting. Secondly, we pay particular attention to the way the subspace is restarted, since classical restarting techniques give a too large subspace in our case. We motivate why it is good enough to just keep the approximation of the searched eigenvector. At the end of the paper we extend this algorithm to shift-and-invert Residual Arnoldi method to calculate the eigenvalue close to a shift σ for a specific parameter dependency. We provide theoretical results and report numerical experiments. The Matlab code is publicly available.

Keywords: Eigenvalues, Tensor-Krylov, Parameter-dependent matrices

2010 MSC: 65F15

1. Introduction

In this paper, we consider parameter-dependent eigenvalue problems

$$\mathbf{A}(\omega)\mathbf{x}(\omega) = \lambda(\omega)\mathbf{x}(\omega), \quad \omega \in \Omega, \quad \mathbf{A} : \Omega \rightarrow \mathbb{R}^{n \times n} : \omega \mapsto \sum_{i=1}^m f_i(\omega)\mathbf{A}_i \quad (1)$$

on a compact parameter set $\Omega \subset \mathbb{R}^d$ and functions $f_i, i = 1, 2, \dots, m$ where m is small relative to n . We look for an eigenpair $(\lambda, \mathbf{x})$ over Ω for which λ is extreme following some ordering, e.g. we sort by modulus if λ is the dominant eigenvalue. We assume that the wanted eigenvalue converges for any ω in a few iterations of the Arnoldi algorithm. In practice, this means
5 that λ should be a well separated extreme eigenvalue of $\mathbf{A}(\omega)$. Here we want to develop a numerical method to efficiently calculate λ for a large set of points in Ω as calculating this eigenvalue for each point individually using standard eigenvalue solvers becomes soon infeasible. At the end of the paper, we extend our method to a case for which the eigenvalue closest to a shift σ is wanted.

If $\mathbf{A}(\omega)$ is symmetric for all $\omega \in \Omega$, a sampling-based subspace method [16] and the Successive Constraint method (SCM)
10 [23] were recently proposed. In the latter, the desired eigenvalue is approximated by iteratively computing lower-and upper

*Corresponding author

Email address: koen.ruymbeek@kuleuven.be (Koen Ruymbeek)

bounds for each parameter value. Another class of solvers are based on polynomial approximation, see for example [1]. This technique is also used for the related problem of calculating statistical moments if the parameter is stochastic, see [3]. The disadvantage of using polynomials is that they suffer from possible lack of smoothness of the desired eigenvalue function [8].

The method that we propose is an example of a Tensor-Krylov method. In these methods, the iteration vectors in the Krylov method are represented by a low rank tensor. Tensor-Krylov techniques are already discussed for various problems, for example to solve a large system of non-linear equations [4, 5] and parametrized linear systems [11].

There are two important aspects in our method. First the iteration vector is represented by a tensor which is truncated to a low-rank representation, for reasons of storage and computational efficiency. Second, the Ritz pairs are obtained from projection of $\mathbf{A}(\omega)$ on the sum of the Krylov spaces for all sample points. A practical issue is the growth of the dimension of this subspace as we add multiple vectors in each iteration. Therefore we often need to restart to keep the algorithm feasible. We gain in efficiency if we need to add less vectors to the subspace (and so introducing a larger error) and have a good way to restart the subspace. We demonstrate that adding fewer vectors in a careful way does not lead to slower convergence and we give an explanation. We give practical bounds for the introduced errors.

The plan of the paper is as follows. In Section 2 we discuss the Arnoldi method and analyse how errors are propagated in the algorithm and give convergence results. We explain the proposed algorithm in the one parameter case and discuss why it is an extension of the Residual Arnoldi algorithm in Section 3. Subsequently, we illustrate the convergence and error propagation with some numerical experiments. In Section 5, we revisit the method for multiple parameters and report a numerical experiment. We extend our method to shift-and-invert Residual Arnoldi when all matrices except one is of rank one and give some numerical examples. We finish this paper with some concluding remarks. The Matlab code can be downloaded following this link.

Throughout the paper, we denote vectors by small bold letters and matrices by capital bold letters. The subspace of $\mathbb{R}^n$ spanned by the columns of an orthonormal matrix $\mathbf{V}$ is denoted by the calligraphic letter $\mathcal{V}$. An index k may be added to refer to the iteration where this subspace is made. With calligraphic letters other than $\mathcal{V}$, we denote tensors of order strictly larger than two. Further, we denote by $\mathbf{I}$ the identity matrix of appropriate dimensions and by the function $\text{diag}(\mathbf{v})$ a diagonal matrix with the elements of the vector $\mathbf{v}$ on its diagonal. With $\|\cdot\|$ we denote the Euclidean norm and we use $\|\cdot\|_F$ to refer to the Frobenius norm.

We end this section by introducing the concept of reduced eigenvalue problem. Let $\mathcal{V} \subset \mathbb{R}^n$ be a subspace and $\mathbf{V}$ an orthonormal matrix where the columns form a basis for $\mathcal{V}$, then we call $\mathbf{V}^T \mathbf{A} \mathbf{V}$ the *reduced eigenvalue problem*. We denote by $(\lambda^\mathcal{V}, \mathbf{x}^\mathcal{V})$ an eigenpair of this reduced eigenvalue problem and the corresponding Ritz vector by $\mathbf{y} = \mathbf{V} \mathbf{x}^\mathcal{V}$. A subindex may be used to indicate the iteration where the subspace is made. We define the corresponding residual as $\mathbf{r} = \mathbf{A} \mathbf{y} - \lambda^\mathcal{V} \mathbf{y}$. We further assume that an eigenvector has norm 1. All proposed algorithms in this paper are implemented in Matlab version R2018a. All experiments are performed on an Intel i5-6300U with 2.5 GHz and 8 GB RAM.

2. Arnoldi method

In this section we introduce the Arnoldi algorithm in a general form, we explain how convergence is measured and give some results regarding its performance when the iterations are contaminated by error. The Arnoldi algorithm itself can be found in pseudo-code in Algorithm 1. In the Arnoldi algorithm, a subspace $\mathcal{V}$ is iteratively built and we look for an approximation of the searched eigenvector in this subspace. In the algorithm, we need to decide about the choice of the continuation vector $\mathbf{t}$, see line 4 in Algorithm 1, which, after multiplying by $\mathbf{A} \mathbf{V}$, is added to the subspace.

Algorithm 1 Arnoldi algorithm

Input: Matrix $\mathbf{A}$, normalized vector $\mathbf{u} \in \mathbb{R}^n$

Output: Approximate eigenpair $(\hat{\lambda}, \mathbf{y})$ of $\mathbf{A}$

```
1:  $\mathbf{w}^\perp = \mathbf{u}$ 
2: for  $k = 1$  to  $n - 1$  do
3:   Add  $\mathbf{w}^\perp$  to the subspace  $\mathcal{V}$ 
4:    $\mathbf{w} = \mathbf{A}\mathbf{V}\mathbf{t}$ 
5:    $\mathbf{w}^\perp = (\mathbf{I} - \mathbf{V}\mathbf{V}^T)\mathbf{w}$ 
6:   Calculate eigenpair  $(\lambda^\mathcal{V}, \mathbf{x}^\mathcal{V})$  of  $\mathbf{V}^T\mathbf{A}\mathbf{V}$ 
7:   Break down if  $(\lambda^\mathcal{V}, \mathbf{x}^\mathcal{V})$  meets a certain convergence criterium
8: end for
9:  $\hat{\lambda} = \lambda^\mathcal{V}$  and Ritz vector  $\mathbf{y} = \mathbf{V}\mathbf{x}^\mathcal{V}$ 
```

It is not difficult to prove that, in exact arithmetic, the eigenvalue approximations do not depend on the used continuation vector as long as the last element in $\mathbf{t}$ is non-zero. This changes when an error (an intended error or a numerical error) is introduced in one of the iterations. In these cases the choice of the continuation vector may be decisive for convergence.

In the original version of the Arnoldi algorithm [2], the k th standard basis vector is used as continuation vector in the k th iteration, but this choice performs badly if there is a significant error in one of the iterations, see [12]. In the next section, another continuation vector is introduced which appears to be more robust against introduced errors. We motivate our choice for the continuation vector and give some convergence results. The convergence behaviour under errors is closely related to inexact Krylov methods, see [22, 21].

In this section we want to answer the question which vector, after multiplying by $\mathbf{A}$, is the best to add to the subspace when this is a Krylov space contaminated with errors. This was the research topic in the papers [28, 27]. In both papers, the authors try to search for the continuation vector $\mathbf{t}$ that maximizes

$$\cos(\theta(\mathbf{x}, \text{Range}\{\mathbf{V}, \mathbf{A}\mathbf{V}\mathbf{t}\})) := \max_{\mathbf{z} \in \text{Range}\{\mathbf{V}, \mathbf{A}\mathbf{V}\mathbf{t}\}} \cos(\theta(\mathbf{x}, \mathbf{z})) \quad (2)$$

The following result is proved in [28].

Theorem 1. *Let $\mathbf{R} = \mathbf{A}\mathbf{V} - \mathbf{V}(\mathbf{V}^T\mathbf{A}\mathbf{V})$, then the continuation vector $\mathbf{t}$ that maximizes (2) is*

$$\mathbf{t} = \mathbf{R}^+\mathbf{x} + \mathbf{b} \quad (3)$$

with $\mathbf{b}$ in the nullspace of $\mathbf{R}$, $\mathbf{R}^+$ the pseudo-inverse of $\mathbf{R}$ and $\mathbf{x}$ the eigenvector associated with λ .

Although this result gives us insight in the choice of the optimal continuation vector, it is not computable as the formula contains the exact eigenvector. The author of [28] argued that in practice, the eigenvector

$$\mathbf{t} = \mathbf{x}^\mathcal{V}. \quad (4)$$

associated with the desired eigenvalue of the reduced problem $\mathbf{V}^T\mathbf{A}\mathbf{V}$ is a good practical choice for the continuation vector. In the remainder of this paper, this is used as continuation vector, despite the fact that in the case of large non-Hermitian matrices, it may be better to use the *refined Ritz vector* [27].

Before the suitability as continuation vector was discussed in [28], it was already observed that this choice leads to a more robust version of the Arnoldi algorithm and was called *Residual Arnoldi* [12, 13]. Throughout this paper, we use the

same name. The name comes from the property that the vector with which we expand the subspace is the residual vector as follows from Lemma 1. The following result is well-known in the literature [19, Proposition 6.8] and is frequently used in the analysis of Krylov methods although it holds for all subspace methods.

Lemma 1. *Let $(\lambda^\mathcal{V}, \mathbf{y})$ be a Ritz vector of $\mathbf{A}$ with respect to the subspace $\mathcal{V}$, then the residual satisfies*

$$\mathbf{r} = (\mathbf{I} - \mathbf{V}\mathbf{V}^T) \mathbf{A}\mathbf{y} \quad (5)$$

so the residual is the result of applying matrix $\mathbf{A}$ on $\mathbf{y}$ orthogonalized with respect to $\mathbf{V}$.

Proof 1. *We can express*

$$\begin{aligned} \mathbf{r} &= \mathbf{A}\mathbf{y} - \lambda^\mathcal{V} \mathbf{y} \\ &= \mathbf{A}\mathbf{V}\mathbf{x}^\mathcal{V} - \mathbf{V}\lambda^\mathcal{V} \mathbf{x}^\mathcal{V} \\ &= \mathbf{A}\mathbf{V}\mathbf{x}^\mathcal{V} - \mathbf{V}(\mathbf{V}^T \mathbf{A}\mathbf{V}) \mathbf{x}^\mathcal{V} \\ &= (\mathbf{I} - \mathbf{V}\mathbf{V}^T) \mathbf{A}\mathbf{V}\mathbf{x}^\mathcal{V}. \end{aligned}$$

Lemma 1 implies that choosing (4) as continuation vector can simplify the Arnoldi algorithm (see Algorithm 1) provided the norm of the residual is used as convergence criterium. The Residual Arnoldi algorithm is shown in Algorithm 2.

Algorithm 2 Residual Arnoldi (RA) Algorithm

Input: Matrix $\mathbf{A}$, initial vector $\mathbf{u} \in \mathbb{R}^n$

Output: Approximate eigenpair $(\hat{\lambda}, \mathbf{y})$ of $\mathbf{A}$

- 1: $\mathbf{r} = \mathbf{A}\mathbf{u} - \lambda^\mathbf{u} \mathbf{u}$, with $\lambda^\mathbf{u}$ the Ritz value associated to the subspace .
 - 2: **for** $k = 1$ to n **do**
 - 3: add $\mathbf{r}$ to the subspace $\mathcal{V}$
 - 4: Calculate eigenpair $(\lambda^\mathcal{V}, \mathbf{x}^\mathcal{V})$ of $\mathbf{V}^T \mathbf{A}\mathbf{V}$
 - 5: $\mathbf{y} = \mathbf{V}\mathbf{x}^\mathcal{V}$
 - 6: Calculate the residual $\mathbf{r}$ of $(\lambda^\mathcal{V}, \mathbf{y})$ via Lemma 1
 - 7: Stop if $\|\mathbf{r}\|$ is small enough
 - 8: **end for**
 - 9: $\hat{\lambda} = \lambda^\mathcal{V}$ and Ritz vector $\mathbf{y} = \mathbf{V}\mathbf{x}^\mathcal{V}$
-

70 **2.1. Convergence of the Residual Arnoldi algorithm**

In this section, we model the error in the RA algorithm, we discuss some convergence results and give tolerances for the different errors such that convergence is assured. We model the possible error $\mathbf{e}$ as an error on $\mathbf{r}$, i.e on the vector that we want to add to the subspace. The error $\mathbf{e}$ consists of two sources: on the one hand, the error $\mathbf{e}_1$ on $\mathbf{x}^\mathcal{V}$ (before the orthogonalization on line 5 in the RA algorithm) and, on the other hand, the error $\mathbf{e}_2$ on $\mathbf{r}$ (after the orthogonalization on line 6). Note that $\mathbf{e}_2$ is already orthogonal to $\mathcal{V}$, because otherwise the added vector is no longer orthogonal to $\mathcal{V}$. To resume, the vector that we add to the subspace can be written as

$$\begin{aligned} \hat{\mathbf{r}} &= (\mathbf{I} - \mathbf{V}\mathbf{V}^T) \mathbf{A}\mathbf{V} (\mathbf{x}^\mathcal{V} + \mathbf{e}_1) + \mathbf{e}_2 \\ &= \mathbf{r} + \underbrace{(\mathbf{I} - \mathbf{V}\mathbf{V}^T) \mathbf{A}\mathbf{V} \mathbf{e}_1}_{\mathbf{e}_1^\perp} + \mathbf{e}_2. \end{aligned} \quad (6)$$

$\underbrace{\hspace{10em}}_{\mathbf{e}}$

The convergence of this method was analysed in [12, 13]. The authors assume that in each iteration k the error $\mathbf{e}^{(k)}$ relative to the residual, is bounded by some η , in the sense

$$\frac{\|\mathbf{e}^{(k)}\|}{\|\mathbf{r}^{(k)}\|} \leq \eta. \quad (7)$$

The main result on which the analysis relies, is that in iteration k there always exists $\mathbf{B}_k$ such that k iterations of the original Arnoldi method on the matrix

$$\hat{\mathbf{A}}_k := \mathbf{A} + \mathbf{B}_k,$$

in the absence of errors results in the same subspace as performing k iterations of RA on matrix $\mathbf{A}$ where in each iteration an error $\mathbf{e}^{(k)}$ is made. In these papers, the authors make the assumption that the size of $\|\mathbf{B}_k\|$ is proportional to η , i.e. there exist C_η such that

$$\|\mathbf{B}_k\| \leq \eta C_\eta.$$

The authors indicate that they cannot prove this assumption, although they have some empirical evidence that it is satisfied. Other assumptions were made on the separation of the spectrum of both $\mathbf{A}$ and $\hat{\mathbf{A}}$.

Theorem 2. *Assume that the above conditions are satisfied, let C_X be a constant depending on C_η and the spectrum of $\mathbf{A}$. If $2\eta C_X < 1$, then*

$$\|\mathbf{y}^{(k)} - \mathbf{x}\| \leq \frac{\kappa_k}{1 - 2\eta C_X}$$

where κ_k is a bound for the difference between the approximation and the searched eigenvector when no error is made. This value goes to zero as k increases.

75

We refer to [13] for a proof and more details. In practice this bound is not (or hardly) computable, so a workable value of η needs to be used in practice.

If the method converges, then the question arises how it depends on η . Empirically, the authors of [12, 13] noticed that the speed of convergence for η lower than or equal to 10^{-3} stays unchanged but that for $1 \gg \eta > 10^{-3}$ the convergence can slow down a bit. For η around 1, we do not have convergence, which can be explained by the fact that the error is as large
80 as the residual.

We now bound the norm of the errors $\mathbf{e}_1$ and $\mathbf{e}_2$ such that convergence is still preserved, i.e. such that condition (7) is fulfilled with $\eta < 10^{-3}$. Let

$$\frac{\|\mathbf{e}_1\|}{\|\mathbf{x}^\vee\|} = \|\mathbf{e}_1\| \leq \eta_X \text{ and } \frac{\|\mathbf{e}_2\|}{\|\mathbf{r} + \mathbf{e}_1^\perp\|} \leq \eta_R \quad (8)$$

be tolerances for the relative error, then we obtain starting from (6)

$$\begin{aligned} \frac{\|\mathbf{e}\|}{\|\mathbf{r}\|} &\leq \frac{\|\mathbf{e}_1^\perp\|}{\|\mathbf{r}\|} + \frac{\|\mathbf{e}_2\|}{\|\mathbf{r}\|} \\ &= \frac{\|(\mathbf{I} - \mathbf{V}\mathbf{V}^T)\mathbf{A}\mathbf{V}\mathbf{e}_1\|}{\|\mathbf{r}\|} + \frac{\|\mathbf{e}_2\|}{\|\mathbf{r}\|} \\ &\leq \underbrace{\frac{\|\mathbf{A}\| \|\mathbf{e}_1\|}{\|\mathbf{r}\|}}_{\leq \eta/2} + \underbrace{\frac{\|\mathbf{e}_2\|}{\|\mathbf{r}\|}}_{\leq \eta/2}. \end{aligned}$$

The aim is now to bound both terms by $\eta/2$. Remark that the first term means that the error propagation due to the matrix-vector product needs to be limited. From (8) it follows that

$$\eta_X \leq \frac{\eta \|\mathbf{r}\|}{2 \|\mathbf{A}\|} \quad (9)$$

which means that η_X scales with the residual norm and the norm of the matrix $\mathbf{A}$. The lower the residual is, the lower η_X needs to be. In the example in Section 4, we demonstrate that if we keep η_X constant, then convergence stagnates if the residual becomes small.

Using $\|\mathbf{r} + \mathbf{e}_1^\perp\| \leq \|\mathbf{r}\| + \|\mathbf{e}_1^\perp\|$, we obtain

$$\begin{aligned} \frac{\|\mathbf{e}_2\|}{\|\mathbf{r}\|} &= \frac{\|\mathbf{e}_2\|}{\|\mathbf{r} + \mathbf{e}_1^\perp\|} \frac{\|\mathbf{r} + \mathbf{e}_1^\perp\|}{\|\mathbf{r}\|} \\ &\leq \frac{\|\mathbf{e}_2\|}{\|\mathbf{r} + \mathbf{e}_1^\perp\|} \frac{\|\mathbf{r}\| + \|\mathbf{e}_1^\perp\|}{\|\mathbf{r}\|} \\ &\leq \frac{\|\mathbf{e}_2\|}{\|\mathbf{r} + \mathbf{e}_1^\perp\|} (1 + \eta/2). \end{aligned}$$

As this needs to be smaller than $\eta/2$, we must impose that

$$\eta_R \leq \frac{\eta}{2 + \eta}. \quad (10)$$

85 3. Tensor-Arnoldi for one parameter

We are ready to propose our algorithm for calculating the wanted eigenvalue over the parameter space. First, a general description of the algorithm is given, then we describe in more detail some specific steps of the algorithm. We explain the algorithm for the case where we have one parameter. For the extension to multiple parameters we refer to Section 5. First, we discretize Ω into n_1 points $\omega_1, \omega_2, \dots, \omega_{n_1}$. We want to compute the desired eigenvalue λ at these points. For this, we
90 consider the linear operator $\mathcal{A}$ such that

$$\mathcal{A}\mathbf{x} = \begin{bmatrix} \mathbf{A}(\omega_1) & & \\ & \ddots & \\ & & \mathbf{A}(\omega_{n_1}) \end{bmatrix} \begin{bmatrix} \mathbf{x}(\omega_1) \\ \mathbf{x}(\omega_2) \\ \vdots \\ \mathbf{x}(\omega_{n_1}) \end{bmatrix} = \begin{bmatrix} \mathbf{A}(\omega_1)\mathbf{x}(\omega_1) \\ \vdots \\ \mathbf{A}(\omega_{n_1})\mathbf{x}(\omega_{n_1}) \end{bmatrix} \in \mathbb{R}^{nn_1}. \quad (11)$$

If $\mathbf{x}$ is reformulated as matrix $\mathbf{X} = [\mathbf{x}(\omega_1), \mathbf{x}(\omega_2), \dots, \mathbf{x}(\omega_{n_1})] \in \mathbb{C}^{n \times n_1}$, then (11) is equivalent to an operator $\mathcal{A}(\mathbf{X})$ such that $\text{vec}(\mathcal{A}(\mathbf{X})) = \mathcal{A}\mathbf{x}$. Conceptually, our method boils down to applying RA to (11) which is nothing more than applying RA for each ω_i individually.

In order to efficiently compute $\mathcal{A}(\mathbf{X})$ we need the expansion in (1). If we have a decomposition $\mathbf{X} = \mathbf{U}\mathbf{Z}$ with $\mathbf{U} \in \mathbb{R}^{n \times n_U}$
95 and $\mathbf{Z} \in \mathbb{C}^{n_U \times n_1}$, then in the same way as in [11] we obtain

$$\mathcal{A}(\mathbf{X}) = \underbrace{[\mathbf{A}_1\mathbf{U}, \mathbf{A}_2\mathbf{U}, \dots, \mathbf{A}_m\mathbf{U}]}_{\mathbf{W}_U \in \mathbb{R}^{n \times n_U m}} \begin{bmatrix} \mathbf{Z} \cdot \text{diag}([f_1(\omega_1), f_1(\omega_2), \dots, f_1(\omega_{n_1})]) \\ \mathbf{Z} \cdot \text{diag}([f_2(\omega_1), f_2(\omega_2), \dots, f_2(\omega_{n_1})]) \\ \vdots \\ \mathbf{Z} \cdot \text{diag}([f_m(\omega_1), f_m(\omega_2), \dots, f_m(\omega_{n_1})]) \end{bmatrix} \underbrace{\quad}_{\mathbf{W}_Z \in \mathbb{C}^{n_U m \times n_1}}. \quad (12)$$

Remark that the term

$$\mathbf{Z} \cdot \text{diag}([f_i(\omega_1), f_i(\omega_2), \dots, f_i(\omega_{n_1})]) = [f_i(\omega_1)\mathbf{Z}(:, 1), f_i(\omega_2)\mathbf{Z}(:, 2), \dots, f_i(\omega_{n_1})\mathbf{Z}(:, n_1)],$$

so column j of $\mathbf{Z}$ is multiplied with the function f_i evaluated in ω_j . Let now $\mathbf{X}^\mathcal{V}$ be the matrix storing in each column i the eigenvector $\mathbf{x}^\mathcal{V}(\omega_i)$ of the reduced eigenvalue problem $\mathbf{V}^T \mathbf{A}(\omega_i) \mathbf{V}$ with respect to the subspace $\mathcal{V}$. We first make a low-rank

decomposition $\mathbf{U}\mathbf{Z}$ of $\mathbf{X}^\mathcal{V}$ and subsequently we compute $\mathbf{W}_U\mathbf{W}_Z = \mathcal{A}((\mathbf{V}\mathbf{U})\mathbf{Z})$ via (12) and we orthogonalise the result with respect to $\mathcal{V}$ to obtain an approximate residual for each parameter value, see Lemma 1. This is analogous to lines 5 and 6 in the RA algorithm (see Algorithm 2). Instead of adding all column vectors of the residual, we make a low-rank approximation and add in the next iteration the resulting basis $\mathbf{U}$ to $\mathcal{V}$. The algorithm for the low-rank decomposition is the same as in [11].

After calculating new eigenvalues and eigenvectors we can repeat the same procedure until all residuals are below a certain tolerance ϵ . As we are adding multiple vectors to the subspace in each iteration and the maximal dimension of the subspace is assumed to be prescribed, the question arises how we restart the subspace when it attains its maximal dimension. In the literature, several techniques are developed for the Arnoldi algorithm like polynomial restart methods [17, 18] or the implicitly restarted method [24]. All of them try to delete the part of the subspace that corresponds to an unwanted part of the spectrum, but still keeps a significant part of the subspace. The reason is that not only the Ritz vector associated with the wanted eigenvalue, but also some other Ritz vectors whose eigenvalues are close to $\lambda^\mathcal{V}$ help improving the convergence to the wanted eigenvalue. This is known as *thick restarting*. In our case of parametric eigenvalue problems, it is slightly more complicated as eigenvectors vary over the parameter. This means that for just representing the wanted eigenvalue we already need quite a large subspace. If also eigenvectors associated with other eigenvalues need to be present in the restarted subspace, there is less room for iterations within the same subspace. Therefore, we choose to just use a basis for the first eigenvectors as restarted subspace. This is a good strategy since it is known from [16, 20] that the derivative of an eigenvector, which is a measure for the change of the eigenvector, depends the most on the eigenvectors whose eigenvalues are the closest to $\lambda(\omega)$. As we estimate the eigenvector over the parameter space, we therefore also estimate the eigenvector whose eigenvalues are the closest to $\lambda(\omega)$. The algorithm can be found in Algorithm 3.

For the calculation of $\mathbf{V}^T \mathbf{A}(\omega_i) \mathbf{V}$, $i = 1, \dots, n_1$, we first calculate $\mathbf{V}^T \mathbf{A}_i \mathbf{V}$, $i = 1, \dots, m$ and sum these smaller matrices appropriately. The operation $\mathbf{A}_i \mathbf{V}$, $i = 1, \dots, m$ is not expensive as these $\mathbf{A}_i$ is usually sparse, but the inner product $\mathbf{V}^T (\mathbf{A}_i \mathbf{V})$ need $2n^2 n_v^2$ operations with n_v the dimension of $\mathbf{V}$. As $\mathbf{V}$ is built iteratively, results from the previous iteration can be used if the subspace was not restarted. The lower η_X and η_R , the higher the number of these inner products. In the example in the Section 4, we count the number of these inner products and compare them for different values of η_X and η_R .

4. Error analysis

We illustrate the performance of the proposed algorithm through an example. We show here what the impact is of the choice of the tolerance levels η_R resp. η_X in the PRA algorithm and of the maximal size of the subspace $\mathcal{V}$. In particular, we explain why choosing η_R relatively high gives the best results, i.e we need less iterations and less expensive matrix-vector products without affecting the convergence.

Example 1. Consider the convection-diffusion differential equation

$$f = c_1 u_x + c_2 u_y + d_{11} u_{xx} + 2d_{12} u_{xy} + d_{22} u_{yy} \quad (13)$$

for $(x, y) \in \Xi = (0, 1) \times (0, 1)$. The real function u has the argument (x, y) . The constants c_1 and c_2 are associated with the convection part and the constants d_{11} , d_{12} and d_{22} are associated with the diffusion part of the equation. It is always assumed that matrix

$$\mathbf{D} = \begin{bmatrix} d_{11} & d_{12} \\ d_{12} & d_{22} \end{bmatrix}$$

Algorithm 3 Parametric Residual Arnoldi (PRA) Algorithm

Input: $\mathbf{A}(\omega)$, orthonormal matrix $\mathbf{U} \in \mathbb{R}^{n \times p}$, tolerance η , tolerances η_X and η_R for the two decompositions, maximal dimension of subspace $n^\mathcal{V}$

Output: Approximate eigenpair $(\hat{\lambda}(\omega_i), \mathbf{y}(\omega_i)), i = 1, 2, \dots, n_1$

```
1:  $\hat{\mathbf{R}} = [1, 1, \dots, 1], \quad k = 1$ 
2: while  $\|\hat{\mathbf{R}}\|_F / \sqrt{n_1} > \epsilon$  do
3:   Add  $\mathbf{U}$  to the subspace  $\mathcal{V}$ 
4:   Calculate eigenpair  $(\lambda^\mathcal{V}(\omega_i), \mathbf{x}^\mathcal{V}(\omega_i))$  of  $\mathbf{V}^T \mathbf{A}(\omega_i) \mathbf{V}, i = 1, 2, \dots, n_1$ 
5:    $\mathbf{X}^\mathcal{V} = [\mathbf{x}^\mathcal{V}(\omega_1), \mathbf{x}^\mathcal{V}(\omega_2), \dots, \mathbf{x}^\mathcal{V}(\omega_{n_1})]$ 
6:   if dimension  $\mathcal{V} > n^\mathcal{V}$  then
7:      $[\mathbf{U}, \mathbf{Z}] = \text{lowrank}(\mathbf{X}^\mathcal{V}, 10^{-12})$ 
8:     Make  $\mathcal{V}$  empty
9:   else
10:     $[\mathbf{U}, \mathbf{Z}] = \text{lowrank}(\mathbf{X}^\mathcal{V}, \eta_X)$ 
11:     $[\mathbf{W}_U, \mathbf{W}_Z] = \mathcal{A}((\mathbf{V}\mathbf{U}) \mathbf{Z})$  using (12)
12:     $\mathbf{W}_U^\perp = (\mathbf{I} - \mathbf{V}\mathbf{V}^T) \mathbf{W}_U$  ( Orthogonalize  $\mathbf{W}_U$  w.r.t  $\mathbf{V}$ )
13:    Residual  $\hat{\mathbf{R}} := [\hat{\mathbf{r}}(\omega_1), \hat{\mathbf{r}}(\omega_2), \dots, \hat{\mathbf{r}}(\omega_{n_1})] = \mathbf{W}_U^\perp \mathbf{W}_Z$ 
14:     $[\mathbf{U}, \sim] = \text{lowrank}(\mathbf{W}_U^\perp, \mathbf{W}_Z, \eta_R);$ 
15:  end if
16:   $k = k + 1$ 
17: end while
18:  $\hat{\lambda}(\omega_i) = \lambda^\mathcal{V}(\omega_i)$  and Ritz vector  $\mathbf{y}(\omega_i) = \mathbf{V}\mathbf{X}^\mathcal{V}(\omega_i), i = 1, 2, \dots, n_1$ 
```

is positive semi-definite. For the discretisation of Ξ we use second-order central finite difference formulas and we assume Dirichlet boundary conditions. In this way, we obtain a matrix $\mathbf{A}$ depending on the values $c_1, c_2, d_{11}, d_{12}, d_{22}$ such that

$$\mathbf{A}(c_1, c_2, d_{11}, d_{22}, d_{12}) = c_1 \hat{\mathbf{A}}_1 + c_2 \hat{\mathbf{A}}_2 + d_{11} \hat{\mathbf{A}}_3 + d_{22} \hat{\mathbf{A}}_4 + d_{12} \hat{\mathbf{A}}_5. \quad (14)$$

For this first example we take $c_2 = 1, d_{11} = 1.1, d_{22} = 1, d_{12} = 1$ and we let c_1 vary in the interval $\Omega = [-2.5, 2.5]$. We discretize this interval in 100 points. This means we can rewrite (13) as

$$\mathbf{A}(c_1) = \mathbf{A}_1 + c_1 \mathbf{A}_2.$$

The desired eigenvalue is the eigenvalue with largest real part, in this case the same as the eigenvalue closest to zero. The region Ξ is discretised into 100 points in each dimension, so, the resulting matrices $\mathbf{A}_i$ are of size $10^4 \times 10^4$. The results can be found in Figure 1 and Table 1. Remark that we need many iterations but tests with individual parameters values reveal that also the Arnoldi algorithm needs many iterations.

In Section 2.1 we observed that, in the RA algorithm, we add the residual in each iteration to the subspace and conclude that this is more robust than classical Arnoldi. In the PRA algorithm, however, the subspace is enriched with an approximate residual for all sample points by making low-rank approximations. For this, we rely on the error analysis made in Section 2.1. At line 10, on the one hand, and at line 14 of PRA algorithm on the other, we make low-rank approximations of the matrix consisting of the eigenvector resp. the approximate residual and so introducing errors $\mathbf{E}_1 := [\mathbf{e}_1(\omega_1), \mathbf{e}_1(\omega_2), \dots, \mathbf{e}_1(\omega_{n_1})]$ and

$\mathbf{E}_2 := [\mathbf{e}_2(\omega_1), \mathbf{e}_2(\omega_2), \dots, \mathbf{e}_2(\omega_{n_1})]$. We first discuss the error $\mathbf{E}_1$. The low-rank approximation is such that the relative error is bounded by the tolerance η_X , i.e

$$\frac{\|\mathbf{E}_1\|_F}{\|\mathbf{X}^\vee\|_F} \leq \eta_X. \quad (15)$$

Since $\|\mathbf{X}^\vee\|_F = \sqrt{n_1}$, it follows that $\|\mathbf{E}_1\|_F \leq \eta_X \sqrt{n_1}$. Technically, we can only prove from this that

$$\|\mathbf{e}_1(\omega_i)\|^2 \leq \eta_X^2 n_1, i = 1, 2, \dots, n_1,$$

but assuming the error is uniformly distributed over the sample points, we can assume that

$$\|\mathbf{e}_1(\omega_i)\| \lesssim \eta_X. \quad (16)$$

Figure 1: Results of the PRA algorithm for different tolerances η_X and η_R with maximal subspace dimension 150 for Example 1. In (a) we plot the Frobenius norm of the residual, scaled with the number of sample points. In (b) the average first error $\mathbf{e}_1$ relative to the average residual, scaled with the norm of $\mathbf{A}_0$ is shown. As long as the value is below or around 10^{-3} , the algorithm converges as if there were no error, see the criterium in (9). If the value is large (> 1) convergence is lost.

Figure 1 (a) shows the average residual over the parameter and we observe that for $\eta_X = 10^{-3}$ the convergence slows down after 400 iterations. This is related to the average relative error $\mathbf{e}_1$ plotted in Figure 1 (b), as this becomes too high such that condition (9) is not fulfilled. Analogously to (16) we can prove that

$$\frac{\|\mathbf{e}_2(\omega_i)\|}{\|\mathbf{r}(\omega_i) + \mathbf{e}_1^\perp\|} \lesssim \eta_R, i = 1, 2, \dots, n_1.$$

As seen in (10), we only need to make sure that $\eta_R \leq \frac{\eta}{2 + \eta}$, where η can be chosen as high as 10^{-3} . This explains why in

Figure 1 (a) choosing $\eta_R = \frac{10^{-3}}{2 + 10^{-3}}$ does not harm convergence. Furthermore, it seems even that convergence is faster for higher η_R . This phenomenon is intriguing and cannot be explained only by a relative error plot. It is generally known in

135

Arnoldi-type algorithms that, with each restart, convergence can slow down. The higher the tolerance η_R the less vectors are added to the subspace, so we do not need to restart so often, hence we can do more iterations within the same subspace and convergence is faster.

In Table 1, we see statistics about the computational cost. As expected, we observe that the number of restarts increases as the tolerances decrease and that the average size of the subspace is independent of the chosen tolerances. As indication of the increased computational cost when taking the tolerances smaller, we count the number of inner products when creating the reduced eigenvalue problem (line 4 in the PRA algorithm). A more detailed explanation on this was given at the end of Section 3.

	$\eta_X = 10^{-3}$ $\eta_R = \frac{10^{-3}}{2 + 10^{-3}}$	$\eta_X = 10^{-12}$ $\eta_R = \frac{10^{-3}}{2 + 10^{-3}}$	$\eta_X = 10^{-3}$ $\eta_R = 5 \times 10^{-13}$	$\eta_X = 10^{-12}$ $\eta_R = 5 \times 10^{-13}$
Number of restarts	7	16	15	38
Avg. size of subspace	79.5233	78.7383	79.545	81.0283
Avg. number of inner products for reduced eigenvalue problem	691.74	1429.9167	1272.9133	3434.36

Table 1: Statistics associated to the results shown in Figure 1.

In the experiments in Figure 1, we observed that if η_X is chosen high, this harms the convergence when the norm of the error $\mathbf{e}_1$ starts to be high relative to the residual. The solution for this is to choose the tolerance η_X in iteration k relative to the residual in the previous iteration. In Figure 2, we compare results when holding η_X constant and when we adapt its value relatively with respect to the residual. We see that we obtain similar results. This is explained by Figure 2(b) as we see that the error $\mathbf{e}_1$ relative to the residual remains constant.

Figure 2: Results of the PRA algorithm for different tolerances η_X and η_R with maximal subspace dimension 150 for Example 1. We compare the results when taking a constant low tolerance η_X with the case that η_X is relative to the residual. In (a) we plot the Frobenius norm of the residual, scaled with the number of sample. We see that in all four cases the algorithm converges. In (b), the average error $\mathbf{e}_1$ relative to the average residual, scaled with the norm of $\mathbf{A}_1$ is shown. As expected, we see that this error relative to the residual stays constant.

The conclusion is that we obtain the best results when taking η_R high and taking η_X relative to the residual. In this way the norm of both errors $\mathbf{e}_1$ and $\mathbf{e}_2$ stay constant with respect to the residual, which preserves convergence if this constant is

low enough ($< 10^{-3}$), see Section 4.

5. More than one parameter

In this section, we extend the PRA algorithm to the case of more than one parameter. Not only all computations need to have an equivalent tensor counterpart, we also need a suitable representation of the tensors as the storage requirement grows exponentially with the dimension. In this section, we first introduce some notation, then we discuss the chosen tensor-format and explain how all operations are adapted.

We assume that Ω is a Cartesian product of intervals in $\mathbb{R}$, so

$$\Omega = \Omega_1 \times \Omega_2 \times \dots \times \Omega_d = [a^1, b^1] \times [a^2, b^2] \times \dots \times [a^d, b^d].$$

Each dimension $i, i = 1, 2, \dots, d$ is discretised into n_i points, making a total of $n_0 = n_1 n_2 \dots n_d$ sample points where we want to approximate the desired eigenvalue. A point ω on this grid is denoted as $(\omega_{i_1}, \omega_{i_2}, \dots, \omega_{i_d})$ for $i_j = 1, 2, \dots, n_j, j = 1, 2, \dots, d$. All tensor equivalents are the calligraphic version of the letters used in the one parameter case. For example, the eigenvectors are stored in tensor $\mathcal{X} \in \mathbb{C}^{n \times n_1 \times \dots \times n_d}$ where

$$\mathcal{X}(:, i_1, i_2, \dots, i_d) = \mathbf{x}(\omega_{i_1}, \omega_{i_2}, \dots, \omega_{i_d}).$$

We define tensor $\mathcal{F} \in \mathbb{R}^{m \times n_1 \times n_2 \times \dots \times n_d}$ as the tensor with all function values, so

$$\mathcal{F}(i_0, i_1, \dots, i_d) = f_{i_0}(\omega_{i_1}, \omega_{i_2}, \dots, \omega_{i_d})$$

for a point $\omega = (\omega_{i_1}, \omega_{i_2}, \dots, \omega_{i_d})$. If $d = 1$, this is the matrix

$$\mathbf{F} = \begin{bmatrix} f_1(\omega_1) & f_1(\omega_2) & \dots & f_1(\omega_{n_1}) \\ f_2(\omega_1) & f_2(\omega_2) & \dots & f_2(\omega_{n_1}) \\ \vdots & & & \\ f_m(\omega_1) & f_m(\omega_2) & \dots & f_m(\omega_{n_1}) \end{bmatrix}.$$

5.1. Tensors

We first state some basic operations that can be performed on tensors, then we describe how to adapt the PRA Algorithm. A detailed overview on tensors can be found in [10].

Definition 1. 1. The mode- i matricization or unfolding of a tensor $\mathcal{Z} \in \mathbb{C}^{n \times n_1 \times \dots \times n_d}$ is a matrix $\mathbf{Z}_{(i)} \in \mathbb{C}^{n_i \times (nn_1 \dots n_{i-1} n_{i+1} \dots n_d)}$ where the columns are the mode- i vectors. The order in which they are placed in $\mathbf{Z}_{(i)}$ is chosen to be canonical.

2. The mode- i multiplication of a matrix $\mathbf{M} \in \mathbb{R}^{k \times n_i}$ with tensor $\mathcal{Z}$ is defined as

$$\mathcal{Y} = (\mathbf{I}, \dots, \mathbf{I}, \mathbf{M}, \mathbf{I}, \dots, \mathbf{I}) . \mathcal{Z} \Leftrightarrow \mathbf{Y}_{(i)} = \mathbf{M} \mathbf{Z}_{(i)}$$

where $\mathbf{M}$ is at the $i + 1$ th position in the tuple.

3. The Hadamard product between two tensors $\mathcal{U}$ and $\mathcal{W}$ of the same dimension, is obtained by the elementwise multiplication and is denoted by $\mathcal{U} ._H \mathcal{W}$.

If $\mathbf{U}$ is the column basis for the tensor $\mathcal{X}$, then we can write

$$\mathcal{X} := (\mathbf{U}, \mathbf{I}, \mathbf{I}, \dots, \mathbf{I}) . \mathcal{Z} \tag{17}$$

with $\mathbf{U} \in \mathbb{R}^{n \times k}$ and $\mathcal{Z} = \mathbb{R}^{k \times n_1 \times n_2 \times \dots \times n_d}$, where $\mathbf{U}$ is a basis for $\mathbf{X}_{(0)}$.

We look now for a suitable tensor representation such that all the computations in the iteration can be done in an efficient way. A strong requirement for the decomposition of $\mathcal{X}^\nu$ besides the storage requirement is that the following operations in PRA Algorithm can be computed efficiently within this decomposition:

1. Calculating a real basis for the column space, even if the tensor contains complex numbers (line 7, 10 and 14 in the PRA Algorithm).
2. Approximating a tensor up to a given tolerance (line 7, 10 and 14), in the same sense as in (15).
3. Computing a decomposition of $\mathcal{W}_Z$ given the decomposition of $\mathcal{Z}$ (line 11).

The Tensor-Train format (TT-format) satisfies all the above conditions. Before we illustrate this we briefly recall the TT-format. A more detailed overview can be found in [15]. A tensor $\mathcal{Y} \in \mathbb{C}^{n_0 \times n_1 \times \dots \times n_d}$ is in TT-format if

$$\mathcal{Y}(i_0, i_1, \dots, i_d) = \mathbf{G}_0^Y(i_0) \mathbf{G}_1^Y(i_1) \dots \mathbf{G}_d^Y(i_d) \quad (18)$$

with $\mathbf{G}_j^Y(i_j)$ a $r_{j-1}^Y \times r_j^Y$ matrix. The length r_{j-1}^Y and width r_j^Y of these matrices $r_i^Y, i = -1, 0, \dots, d$ are called the *TT-ranks*. It can immediately be observed that $r_{-1}^Y = r_d^Y = 1$. We call $r_i^Y, i = 0, 1, \dots, d-1$ the *interior* ranks. The storage of the TT-representation is lower than $(d+1)nr^2$ with $r_i^Y \leq r$ and $n_i \leq n$, which is for r small significantly smaller than the storage of the full tensor. The matrices $\mathbf{G}_i^Y$ can be stored in a tensor $\mathcal{G}_j^Y \in \mathbb{R}^{r_{j-1}^Y \times n_j \times r_j^Y}, j = 0, 1, \dots, d$ as $\mathcal{G}_j^Y(:, i_j, :) = \mathbf{G}_j^Y(i_j), j = 1, 2, \dots, n_j$.

For computing a TT-decomposition of a tensor, we perform for each mode consecutively a matricization, an SVD and a reshape, see [15, Algorithm 1]. Approximating a tensor for a given tolerance in the same sense as (15) can also be done in the TT-decomposition, see [15, Theorem 2.2].

The last requirement was that we need to be able to calculate a TT-decomposition of $\mathcal{W}_Z$ (the equivalent of $\mathbf{W}_Z$ in (12)), directly from the decomposition of $\mathcal{Z}$. To see this, we recall the case of $d = 1$. Let $\mathbf{1}_q$ be a vector with all ones of length q .

We first note that

$$\mathbf{Z} \cdot \text{diag}([f_i(\omega_1), f_i(\omega_2), \dots, f_i(\omega_{n_1})]) = \mathbf{Z} \cdot \mathbf{H}(\mathbf{1}_{n_U} [f_i(\omega_1), f_i(\omega_2), \dots, f_i(\omega_{n_1})])$$

from which it follows that

$$\mathbf{W}_Z = ((\mathbf{1}_{n_U} \otimes \mathbf{I}_m) \mathbf{F}) \cdot \mathbf{H}((\mathbf{I}_{n_U} \otimes \mathbf{1}_m) \mathbf{Z}).$$

Note that we use here the left-kronecker product $\otimes$. This means that in case of multiple parameters, $\mathcal{W}_Z$ can be generalised to

$$\mathcal{W}_Z = ((\mathbf{1}_{n_U} \otimes \mathbf{I}_m), \mathbf{I}, \dots, \mathbf{I}) \cdot \mathcal{F} \cdot \mathbf{H}((\mathbf{I}_{n_U} \otimes \mathbf{1}_m, \mathbf{I}, \dots, \mathbf{I}) \cdot \mathcal{Z}.$$

From [15], we obtain that

$$\begin{cases} \mathbf{G}_j^{W_Z}(i_j) &= \mathbf{G}_j^Z(i_j) \otimes \mathbf{G}_j^F(i_j) & j = 1, 2, \dots, d, i_j = 1, 2, \dots, n_j, \\ \mathbf{G}_0^{W_Z} &= \mathbf{G}_0^Z \otimes \mathbf{G}_0^F & j = 0. \end{cases} \quad (19)$$

Note that the TT-decomposition is not unique. Moreover, it is possible that the ranks of the decomposition in (19) of the tensor $\mathcal{W}_Z$ are overestimated.

If all functions $f_i, i = 1, 2, \dots, m$ are separable, then there is no need to first calculate a TT-decomposition of $\mathcal{F}$. We can get a decomposition directly from the separability of the functions f_i . As the functions f_i are separable there are functions

$g_{i_j}^i(\omega_{i_j})$ such that $f_i(\omega_{i_1}, \omega_{i_2}, \dots, \omega_{i_d}) = g_{i_1}^i(\omega_{i_1})g_{i_2}^i(\omega_{i_2}) \dots g_{i_d}^i(\omega_{i_d})$. It can be seen that a possible TT-decomposition is of the form

$$\begin{cases} \mathbf{G}_0^F &= \mathbf{I}_m \\ \mathbf{G}_j^F(i_j) &= \text{diag}([g_1^j(\omega_{i_j}), g_2^j(\omega_{i_j}), \dots, g_m^j(\omega_{i_j})]) \quad j = 1, 2, \dots, d-1, i_j = 1, 2, \dots, n_j, \\ \mathbf{G}_d^F(i_d) &= [g_1^d(\omega_{i_d}), g_2^d(\omega_{i_d}), \dots, g_m^d(\omega_{i_d})]^T \quad i_d = 1, 2, \dots, n_d. \end{cases} \quad (20)$$

In this case, all interior ranks are m . This means that the interior ranks of the tensor $\mathcal{W}_Z$ are maximal $mr_i^Z, i = 0, 1, \dots, d$.

185 Note that also here the ranks of the decomposition are possibly overestimated. The lower are the ranks of the tensor $\mathcal{Z}$, the more efficient the other operations are. To have this, it is important how the eigenvectors are stored in the tensor $\mathcal{X}^\nu$. Eigenvectors are only unique up to a complex constant, but for the ranks of the tensor $\mathcal{X}^\nu$ it makes a difference which constant is used. The best we can do is to compare the eigenvector $\mathbf{x}(\omega)$ for a point ω with a certain vector $\mathbf{a}$, and store instead of $\mathbf{x}(\omega)$ the vector $\frac{(\mathbf{a}^H \mathbf{x}(\omega))}{\|\mathbf{a}^H \mathbf{x}(\omega)\|} \mathbf{x}(\omega)$. Remark that if $\mathbf{a}$ and the eigenvector have no complex elements then this boils
190 down to multiplying with 1 or -1 . A possible choice for $\mathbf{a}$ is to take the eigenvector corresponding to a sample point.

If there are not many points in each dimension, it might be better not to use tensors but to reshape the tensor $\mathcal{X}^\nu$ as a matrix and to apply the PRA algorithm as described in Algorithm 3. The advantage of using a decomposition of the tensor is to reduce the storage and to reduce the number of operations, but if the number of parameter values is low in each dimension, then we cannot gain from the TT-representation as the cost of decomposing is too high in comparison with the
195 possible gain in the remaining steps of the algorithm.

5.2. Numerical Experiments

In this section we extend Example 1 to multiple parameters and give the convergence results. The code can be downloaded following this link.

Example 2. Consider the convection-diffusion differential equation (13) in Example 1. We consider the case where $d_{12} = 0$, we vary the variables c_1 and c_2 both in the interval $[4, 6]$ and we let the variable $d := d_{11} = d_{22}$ vary in the interval $[-1.1, -0.9]$. In this way we can rewrite (14) into the form

$$\mathbf{A}(c_1, c_2, d) = c_1 \mathbf{A}_1 + c_2 \mathbf{A}_2 + d \mathbf{A}_3.$$

The desired eigenvalue is the eigenvalue with smallest real part, in this case the same as the eigenvalue closest to zero. The
200 matrices are of size $10^3 \times 10^3$ and we choose $n_1 = n_2 = n_3 = 30$ resulting in a total of $n_0 = 27 \times 10^3$ parameter values. As tolerances we choose $\eta_R = \frac{10^{-3}}{2 + 10^{-3}}$ and $\eta_X^{(k)} = 10^{-3} \left\| \hat{\mathbf{R}}^{(k)} \right\|_F / \sqrt{n_0}$

We first solve the problem discretised in each dimension by 10 interior discretisation points. We see the result in Figure 3 (a). We observe that we need many iterations for convergence to the solution which is not surprising as using Arnoldi's method for all parameter points separately needs approximately the same number of iterations. The obtained subspace $\mathcal{V}$ is then
205 used as initial subspace when discretising each dimension in 30. We see in Figure 3 (b) that we only need a few iterations.

6. Parametric shift-inverted Residual Arnoldi method

in each iteration of the residual Arnoldi method, we applied the operator $\mathcal{A}$ (see (11)) or its tensor counterpart to Ritz vectors contaminated with error. In this way we are able to compute eigenvalues which are extreme in some sense. Users are often interested in the eigenvalue close to a certain shift σ . For this, the shift-inverted Arnoldi method can be used. For the
210 ease of notation we assume that $\sigma = 0$.

Figure 3: Results of the PRA algorithm with maximal subspace dimension 100 for Example 2. In (a), we see the convergence when taking $n_1 = n_2 = n_3 = 10$. We see that we need quite a large number of iterations but tests showed us that using Arnoldi for all parameter points separately needs approximately the same number of iterations. If we take the resulting subspace $\mathcal{V}$ and we use this as the initial subspace for $n_1 = n_2 = n_3 = 30$, we see that we just need a few iterations

In this case we need to be able to efficiently compute linear systems with $\mathbf{A}(\omega)$ for all parameter samples. The inverse of $\mathbf{A}(\omega)$ can in general not be written in the form (1). However if all matrices except one is of rank one, i.e $\mathbf{A}(\omega) = f_0(\omega)\mathbf{A}_0 + \sum_{i=1}^m f_i(\omega)\mathbf{A}_i$ with $\mathbf{A}_i = \mathbf{u}_i\mathbf{q}_i^T, i = 1, \dots, m$ with m a small number, it can be proven via the Sherman-Morrison-Woodbury formula [9, 26] that for a vector $\mathbf{q}$, it holds that

$$\mathbf{A}(\omega)^{-1}\mathbf{q} \in \text{span}(\mathbf{A}_0^{-1}\mathbf{q}, \mathbf{A}_0^{-1}\mathbf{u}_1, \mathbf{A}_0^{-1}\mathbf{u}_2, \dots, \mathbf{A}_0^{-1}\mathbf{u}_m). \quad (21)$$

We see that there is only one basis vector of the subspace that depends on $\mathbf{q}$, while other vectors are the same for any $\mathbf{q}$ and for all parameter samples. Hence, one could keep all other vectors during all iterations in subspace $\mathcal{V}$.

It is possible to directly apply the PRA algorithm on $\mathbf{A}(\omega)^{-1}$, however, in [13, 12], the authors proposed an adaptation of RA. In exact arithmetic, if we apply the RA algorithm on $\mathbf{A}^{-1}$, we add the residual vector $\mathbf{r}_{A^{-1}} = \mathbf{A}^{-1}\mathbf{y} - \frac{1}{\lambda}\mathbf{y}$ with $(\lambda, \mathbf{y})$ a Ritz-pair of $\mathbf{A}$. The authors proposed to first calculate the residual $\mathbf{r} = \mathbf{A}\mathbf{y} - \lambda\mathbf{y}$, to solve the system

$$\mathbf{A}\mathbf{u} = \mathbf{r}. \quad (22)$$

and to add the resulting $\mathbf{u}$ to the subspace. The authors in [13, 12] proved that $\mathbf{u}$ is parallel to $\mathbf{r}_{A^{-1}}$. This has the advantage that solving (22) inexactly, produces an error proportional to $\mathbf{r}$.

As the authors of [13, 12] claim that also here the error proportional to the residual can be as high as 10^{-3} without harming the convergence, it is more appropriate to calculate the vector via (22). Their method was called the Shift-and-invert Residual Arnoldi algorithm and similarly to the PRA algorithm, we extend this method to the parametric case. Instead of (22) we only need to calculate

$$\mathbf{A}_0\mathbf{u}(\omega) = \mathbf{r}(\omega)$$

upto a relative tolerance η_R , as we keep the other basis vectors in (21) in the subspace $\mathcal{V}$ in all iterations. The tolerance for the relative error when solving (22) is denoted with η_R and the tolerance for the low-rank approximation of the eigenvector over the parameter, is denoted with η_X .

6.1. Numerical experiments

The following examples arise from spectrum-based stabilization methods for time-delay systems [7, 14]. Time-delay stability can be inferred from the spectral abscissa (real part of the rightmost eigenvalue) of an eigenvalue problem, discretizing the infinitesimal generator operator by collocation approach [6]. Spectrum-based stabilization is based on minimizing the spectral abscissa as a function of controller parameter and favors situations where the spectral abscissa is not smooth, since minima of the spectral abscissa are often characterized by defective rightmost eigenvalues [8]. In this context, our novel approach may lead to a significant contribution in the probabilistic spectrum-based stabilization [7] since it permits to tackle larger scale problem with higher stochastic dimension. The code can be downloaded following this link.

Example 3. The spectral abscissa function, associated to the oscillator with feedback delay system [8, Example 2], is approximated by the Infinitesimal Generator approach [6] as the spectral abscissa of the following parametric matrix:

$$\mathbf{A}(\omega_1, \omega_2; K_1, K_2) = \mathbf{A}_0(K_1, K_2) - \omega_1^2 \mathbf{A}_1 + -2\omega_1\omega_2 \mathbf{A}_2, \text{ with } \omega_1 \in [0.9, 1.1], \omega_2 \in [0.1, 0.2]. \quad (23)$$

The non-singular matrix $\mathbf{A}_0$ and the rank-1 matrices $\mathbf{A}_1, \mathbf{A}_2$ have dimension 1002×1002 and are given by:

$$\mathbf{A}_0(K_1, K_2) = \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & 0 & \cdots & 0 & K_1 & K_2 \\ & & & \mathbf{I}_2 \otimes \mathbf{D}_M & & & \end{pmatrix}, \mathbf{A}_1 = \begin{pmatrix} 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix} \begin{pmatrix} 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}^T, \mathbf{A}_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}^T$$

where $\mathbf{D}_M \in \mathbb{R}^{M \times M}$ is the differentiation matrix [25, Chapter 6] obtained with $M = 500$ Chebyshev nodes as collocation points. The controller parameters K_1 and K_2 are selected so that the spectral abscissa function presents different behaviors on the parametric domain $(\omega_1, \omega_2) \in [0.9, 1.1] \times [0.1, 0.2]$ as illustrated in Figures 4 (a) and (b): for $(K_1, K_2) = [0.2, 0.2]$ the spectral abscissa smoothly behaves, while for $(K_1, K_2) = [0.6179, -0.0072]$ the spectral abscissa presents non-Lipschitz continuous points, due to multiple non-semi-simple rightmost eigenvalues. We refer to [8] for further information on the example.

We use $\eta_R = \frac{10^{-3}}{2 + 10^{-3}}$ fixed and $\eta_X^{(k)} = 10^{-3} \left\| \hat{\mathbf{R}}^{(k)} \right\|_F / \sqrt{n_0}$ relative to the residual as explained in Section 4. The parameter intervals are both discretised into 30 points, so we have a total of $n_0 = 900$ points for which the spectral abscissa needs to be found. The spectral abscissa is plotted in Figure 4 (a)-(b). The convergence results are shown in Figure 4 (c)-(d).

Example 4. We analyze the spectral abscissa function associated to the heat transfer set up with two static feedback controllers $\mathbf{K}$, designed by probabilistic stabilization method considering the temperature on the left cooler and both cooler affected by uncertainty [7, Section 5.3]. The spectral abscissa function is determined by the eigenvalues of the following 1030×1030 parametric matrix

$$\mathbf{A}(\omega_1, \omega_2; \mathbf{K}) = \mathbf{A}_0(\mathbf{K}) + \frac{1 - \omega_1}{\omega_1} \mathbf{A}_1 + \frac{1 - \omega_2}{\omega_2} \mathbf{A}_2, \text{ where } \omega_1 \in [13.5, 16.5], \omega_2 \in [15.3, 18.7].$$

The two static controllers $\mathbf{K}$, designed in order to minimize the mean of the spectral abscissa, modelling the uncertainty on the temperatures by independent uniform random variables [7, Section 5.3], lead to different behaviors of the spectral abscissa. In Figure 5 (a) resp. (b), the spectral abscissa for $\mathbf{K} = [0.5105, -9.1810 \times 10^{-2}]$ resp. $\mathbf{K} = [0.45043, -0.1631]$ is used. The parameter intervals are both discretised into 30 points, so we have a total of $n_0 = 900$ parameter values for which the spectral abscissa needs to be found. We use $\eta_R = \frac{10^{-3}}{2 + 10^{-3}}$ fixed and $\eta_X^{(k)} = 10^{-3} \left\| \hat{\mathbf{R}}^{(k)} \right\|_F / \sqrt{n_0}$ relative to the residual as explained in Section 4. The convergence results are shown in Figure 5 (c)-(d).

(a)

(b)

(c)

(d)

Figure 4: The spectral abscissa over the parameters for the two different cases in Example 3. We see in (a) that the spectral abscissa varies smoothly over the parameters, in (b) this is not the case anymore. The convergence results for the two different cases in Example 3 are plotted in (c) and (d). We see that in the two cases the algorithm converges within eleven iterations.

(a)

(b)

(c)

(d)

Figure 5: In (a)-(b), The spectral abscissa over the parameters for the two different cases in Example 4. The convergence results for the two different cases in Example 4 can be seen in (c)-(d).

245 We see in Figures 4 (c)-(d) and 5 (c)-(d) that in both examples only a few iterations were needed. In Figure 4 (b) and 5 (a)-(b) it can be observed that the spectral abscissa over the parameter is not differentiable as in some points (we only count eigenvalues with imaginary part equal or greater than zero) there are multiple eigenvalues for which the real part collides. We need to be careful when some eigenvalues of the large eigenvalues are close to each other. In our experience, subspace methods for calculating eigenvalues of non-normal matrices, can fail to converge, if the searched eigenvalue is close to other
250 eigenvalues. Therefore, it is possible that due to the projection an eigenvalue is selected which is not the one we look for. The technique that we apply, is to save for every parameter sample in iteration k five eigenvectors of the projected problem that best satisfies our criterium. In the following iteration $k + 1$ only an eigenvalue with associated eigenvector is selected if the angle between this eigenvector and one of the five selected eigenvectors in iteration k is below a prescribed tolerance.

7. Conclusion

255 In this paper, we proposed a method to calculate eigenvalues of a parametrized matrix $\mathbf{A}(\omega)$ for a number of parameter values at the same time. We started with introducing the Residual Arnoldi algorithm, which is a version of the Arnoldi algorithm that is robust against introduced errors. We applied the Residual Arnoldi Algorithm on a large number of sample points at the same time. To keep the computations feasible, we made low-rank approximations. On the one hand, we made a low-rank approximation of the searched eigenvector of the reduced eigenvalue problem over the parameters with a relative
260 tolerance of η_X , and on the other, we made a low-rank approximation on the residual over the parameters with a relative tolerance of η_R . We saw that the total error relative to the residual needs to be smaller than 10^{-3} , from which we deduce that η_X needs to be chosen proportional to the residual and that η_R can be chosen as high as $\frac{10^{-3}}{2 + 10^{-3}}$, see Section 2.1.

We illustrated with numerical experiments that a higher η_R even improves performance, since a higher tolerance leads to a slower increase of the subspace dimension and consecutely less restarts. At the end of the paper, we made the extension to
265 multiple parameters and to shift-inverted Arnoldi for a specific range of problems. A possible extension is to make a version of the parametric shift-inverted Arnoldi method when more than one matrix has rank higher than one. The difficulty in this problem, is that we need to be able to solve parametrized linear systems efficiently. Another possibility for further research is the calculation of statistical moments of certain eigenvalues over the parameter.

Acknowledgements

270 The authors thank the referees for their helpful remarks. We also thank Nick Vannieuwenhoven for some fruitful discussions about the used tensor format and we thank Luca Fenzi for the examples used in Section 6.1. This work was supported by the project KU Leuven Research Council grant C14/17/072 and by the project G0A5317N of the Research Foundation-Flanders (FWO - Vlaanderen). This work does not have any conflicts of interest.

References

- 275 [1] Andreev, R., Schwab, C., 2012. Sparse Tensor Approximation of Parametric Eigenvalue Problems, in: Graham, I.G., Hou, T.Y., Lakkis, O., Scheichl, R. (Eds.), Numerical Analysis of Multiscale Problems. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 203–241. doi:10.1007/978-3-642-22061-6_7.
- [2] Arnoldi, W.E., 1951. The principle of minimized iterations in the solution of the matrix eigenvalue problem. Quarterly of Applied Mathematics doi:10.1090/qam/42792.

- [3] Benner, P., Onwunta, A., Stoll, M., 2019. A Low-Rank Inexact Newton-Krylov Method for Stochastic Eigenvalue Problems. Computational Methods in Applied Mathematics doi:10.1515/cmam-2018-0030.
- [4] Bouaricha, A., 1996. Tensor-Krylov methods for large nonlinear equations. Computational Optimization and Applications doi:10.1007/BF00248265.
- [5] Bouaricha, A., Schnabel, R.B., 1998. Tensor methods for large sparse systems of nonlinear equations. Mathematical Programming, Series B doi:10.1007/BF01580076.
- [6] Breda, D., Maset, S., Vermiglio, R., 2015. Stability of linear delay differential equations: A numerical approach with MATLAB, in: SpringerBriefs in Control, Automation and Robotics.
- [7] Fenzi, L., Michiels, W., 2017. Robust stability optimization for linear delay systems in a probabilistic framework. Linear Algebra and Its Applications doi:10.1016/j.laa.2017.03.020.
- [8] Fenzi, L., Michiels, W., 2019. Polynomial (chaos) approximation of maximum eigenvalue functions. Numerical Algorithms doi:10.1007/s11075-018-00648-9.
- [9] Hager, W.W., 1989. Updating the inverse of a matrix. SIAM Review doi:10.1137/1031049.
- [10] Kolda, T.G., Bader, B.W., 2009. Tensor Decompositions and Applications. SIAM Review doi:10.1137/07070111X.
- [11] Kressner, D., Tobler, C., 2011. Low-Rank Tensor Krylov Subspace Methods for Parametrized Linear Systems. SIAM Journal on Matrix Analysis and Applications doi:10.1137/100799010.
- [12] Lee, C.R., 2007. Residual Arnoldi Methods : Theory , Package , and Experiments Professor G . W . Stewart Department of Computer Science. Phd thesis. University of Maryland.
- [13] Lee, C.R., Stewart, G., 2007. Analysis of the residual Arnoldi method, Technical Report. Technical Report. University of Maryland.
- [14] Michiels, W., 2011. Spectrum-based stability analysis and stabilisation of systems described by delay differential algebraic equations. IET Control Theory and Applications doi:10.1049/iet-cta.2010.0752.
- [15] Oseledets, I.V., 2011. Tensor-Train Decomposition. SIAM Journal on Scientific Computing doi:10.1137/090752286.
- [16] Ruymbeek, K., Meerbergen, K., Michiels, W., 2019. Calculating the minimal/maximal eigenvalue of symmetric parametrized matrices using projection. Numerical Linear Algebra with Applications 26.
- [17] Saad, Y., 1980. Variations on Arnoldi's method for computing eigenelements of large unsymmetric matrices. Linear Algebra and Its Applications doi:10.1016/0024-3795(80)90169-X.
- [18] Saad, Y., 1984. Chebyshev Acceleration Techniques for Solving Nonsymmetric Eigenvalue Problems. Mathematics of Computation doi:10.2307/2007602.
- [19] Saad, Y., 2011. Numerical Methods for Large Eigenvalue Problems. Numerical Methods for Large Eigenvalue Problems , 285doi:10.1137/1.9781611970739.
- [20] Seyranian, A.P., Mailybaev, A.A., 2003. Multiparameter stability theory with mechanical applications. volume 13. World Scientific.

- [21] Simoncini, V., 2005. Variable Accuracy of Matrix-Vector Products in Projection Methods for Eigencomputation. SIAM Journal on Numerical Analysis doi:10.1137/040605333.
- 315 [22] Simoncini, V., Szyld, D.B., 2003. Theory of Inexact Krylov Subspace Methods and Applications to Scientific Computing. SIAM Journal on Scientific Computing doi:10.1137/s1064827502406415.
- [23] Sirković, P., Kressner, D., 2016. Subspace Acceleration for Large-Scale Parameter-Dependent Hermitian Eigenproblems. SIAM Journal on Matrix Analysis and Applications 37, 695–718. doi:10.1137/15M1017181.
- [24] Sorensen, D.C., 1992. Implicit Application of Polynomial Filters in a k-Step Arnoldi Method . SIAM Journal on Matrix
320 Analysis and Applications doi:10.1137/0613025.
- [25] Trefethen, L.N., 2000. Spectral Methods in MATLAB. doi:10.1137/1.9780898719598.
- [26] Woodbury, M., 1950. Inverting modified matrices. Technical Report.
- [27] Wu, G., Zhang, L., 2014. On expansion of search subspaces for large non-Hermitian eigenproblems. Linear Algebra and Its Applications doi:10.1016/j.laa.2014.04.021.
- 325 [28] Ye, Q., 2008. Optimal expansion of subspaces for eigenvector approximations. Linear Algebra and Its Applications doi:10.1016/j.laa.2007.08.021.