

Backward errors for multiple eigenpairs in structured and unstructured nonlinear eigenvalue problems

Miryam Gnazzo*, Leonardo Robol†

October 28, 2025

Abstract

Given a nonlinear matrix-valued function $F(\lambda)$ and approximate eigenpairs (λ_i, v_i) , we discuss how to determine the smallest perturbation δF such that $[F + \delta F](\lambda_i)v_i = 0$; we call the distance between the F and $F + \delta F$ the *backward error* for this set of approximate eigenpairs. We focus on the case where $F(\lambda)$ is given as a linear combination of scalar functions multiplying matrix coefficients F_i , and the perturbation is done on the matrix coefficients. We provide inexpensive upper bounds, and a way to accurately compute the backward error by means of direct computations or through Riemannian optimization. We also discuss how the backward error can be determined when the F_i have particular structures (such as symmetry, sparsity, or low-rank), and the perturbations are required to preserve them. For special cases (such as for symmetric coefficients), explicit and inexpensive formulas to compute the δF_i are also given.

1 Introduction

We consider matrix-valued functions $F : \mathbb{C} \mapsto \mathbb{C}^{n \times n}$ and the related nonlinear eigenvalue problem, that consists in finding $\lambda \in \mathbb{C}$ and v such that

$$F(\lambda)v = 0, \quad v \in \mathbb{C}^n \setminus \{0\}.$$

The pair (λ, v) is called *eigenpair*, and λ , v are an *eigenvalue* and an *eigenvector* for $F(\lambda)$, respectively. Quite often, the matrix-valued function $F(\lambda)$ is given in *split form* as a linear combination of matrix coefficients multiplied by analytic functions:

$$F(\lambda) = f_1(\lambda)F_1 + \dots + f_k(\lambda)F_k. \quad (1)$$

The coefficients F_j frequently encode data coming from the underlying application (for instance, the coefficients of a stiffness or damping matrix in the PDE setting). A few representative cases of this application can be found in [5]; for instance see the quadratic eigenvalue problem **spring** associated with a damped mass-spring system [21, Example 2]. It is often useful to consider non-polynomial scalar functions f_j , such as exponentials. This kind of matrix-valued functions arise in the context of constant-coefficients delay differential equations, see for instance [8]. We will only consider nonlinear eigenvalue problems of the form (1) in this work.

Nonlinear eigenvalue problems arise in a wide set of applications; a large collection of examples can be found in the MATLAB package **nlevp** [5]. Several algorithms have been developed for

*Division of Mathematics, Gran Sasso Science Institute, L'Aquila, Italy (miryam.gnazzo@gssi.it).

†Dipartimento di Matematica, Università di Pisa, Italy (leonardo.robol@unipi.it).

the numerical solution of nonlinear eigenvalue problems, see for instance [12] for a survey on the nonlinear eigenvalue problem. An implementation of possible solvers for this problem is available in the Julia package NEP-PACK [15]. Moreover, recent solvers employ a variant of the AAA algorithm for the solution of the nonlinear eigenvalue problem [17].

Selected instances of this problem have been thoroughly studied in the literature: when $f_j(\lambda) = \lambda^{j-1}$ we obtain a *polynomial eigenvalue problem*. For instance a complete review on the quadratic eigenvalue problem can be found in [22]. In the case $k = 2$ we get a matrix pencil, or a standard eigenvalue problem.

In numerical linear algebra, the standard way to assess the quality of computed eigenvalues and eigenvectors is to determine their backward error; the latter is defined as the distance from the closest eigenvalue problem for which the computed eigenvalues (and eigenvectors) are exact.

For standard eigenvalue problems, we have explicit formulas that relate the residual norm $\|F(\lambda)v\|_2$ to the backward error, and that are part of any numerical linear algebra textbook; similar results can be given for polynomial eigenvalue problems [21]. Moreover in [9] the authors proposed a backward error analysis for the solution of the polynomial eigenvalue problems and complete polynomial eigenproblems, via block Kronecker linearizations. Some results can be found for more general nonlinear eigenvalue problems as well. In [3], a formulation for the backward error of a given eigenpair has been proposed in the context of homogeneous nonlinear eigenvalue problems, with particular attention to structured matrix-valued functions F . For rational eigenvalue problems, in [19] the authors derive formulae for the symmetric backward error of one eigenvalue. In [16] the authors develop a characterization for the backward error associated with a set of eigenvalues for a matrix-valued analytic function. Their bound however does not relate with the split form of (1), but rather focus on finding a small functional perturbation.

The contribution of this work is twofold:

1. We provide computable and inexpensive bounds for the backward error of a *set of eigenvalues and eigenvectors* (or for just the eigenvalues, if no eigenvectors have been computed).
2. We give numerical procedures based on Riemannian optimization that compute the backward error accurately, and that do so retaining any structure found in the coefficients F_j (such as sparsity, low-rank, symmetries, ...).

We also provide computable bounds for the backward error in the structured case, but these will be inexpensive only for the case of symmetric nonlinear eigenvalue problems. For more general structures, we found that directly computing the backward error is often the best way to proceed.

A previous attempt at characterizing structured backward errors for nonlinear eigenvalue problems can be found in [3], where the authors focus on only one eigenpair. Our bounds will reduce to the one in this work when considering a set with a single eigenpair.

We remark that it may be tempting to use results for a single eigenpair to draw conclusions on the accuracy of a set of eigenpairs, but this can be misleading. Indeed, it may happen that the backward errors of two different approximate eigenvalues λ_1, λ_2 are small, but there is no close-by nonlinear eigenvalue problem that has both as eigenvalues. An example showcasing this possibility may be found in Section 6.2 in [16].

A similar discussion of considering a set of eigenvalues at once for standard eigenvalue problems can be found in [20], together with a complete survey on structured normwise backward errors for a set of eigenpairs. Following [20], we define the backward error associated with a set of p eigenpairs $(\hat{\lambda}_i, \hat{v}_i)$ as follows:

$$\eta := \min \left\{ \|\delta F_1, \dots, \delta F_k\|_F \mid \sum_{j=1}^k f_j(\hat{\lambda}_i)(F_j + \delta F_j)\hat{v}_i = 0, 1 \leq i \leq p \right\}.$$

This measure does not take into account possible additional structures on the coefficient matrices F_j . Nevertheless it may be useful to include structures into account, for instance in situations where we would like to exploit the structure of the problem and therefore the sensitivity of the solution of the nonlinear eigenvalue problem should be measured with respect to the same structure. This is done for instance in [10] where the authors propose a computable structured condition number for the class of parametrized quasiseparable matrices. Results on condition numbers of structured matrix polynomials, such symmetric and palindromic, have been developed by Adhikari et al. in [2]. In addition, the authors analyze the relations with the condition number of structured linearizations of the matrix polynomials. For all contexts where the coefficient belong to prescribed classes of structured matrices, we will also introduce a structured backward error η_S defined analogously, but with the constraint of $F_j + \delta F_j$ sharing the same structure of F_j . We will discuss in detail the cases of an assigned sparsity pattern, a maximum rank, and symmetries. The proposed algorithm will be able to deal with different structures for each coefficient, and also multiple structures at once with no modifications.

The paper is organized as follows. In Section 2, we analyze the unstructured backward error for the nonlinear eigenvalue problem and provide several computable upper bounds for it. In Section 3, we present an overview of structured backward errors for a set of approximated eigenpairs. For the case of general linear structures, we provide a formula for the backward error. Then we specialize the results for the symmetry structure, providing a cheaper computable upper bound for the backward error. In the end, we consider nonlinear structures, such as low-rank ones, for which we are able to provide an upper bound computed through the use of a Riemannian optimization-based technique. In Section 4 a few numerical tests and examples conclude the paper.

2 Backward errors for nonlinear eigenvalue problems

We consider matrix-valued functions $F : \mathbb{C} \mapsto \mathbb{C}^{n \times n}$ in split form (1), that is $F(\lambda) := \sum_{j=1}^k F_j f_j(\lambda)$, where $F_j \in \mathbb{C}^{n \times n}$ and $f_j : \mathbb{C} \mapsto \mathbb{C}$ analytic functions for $j = 1, \dots, k$. Observe that given a general matrix-valued function, it is always possible to write it in split form, decomposing it as $[F(\lambda)]_{ij} e_i e_j^T$, for $i, j = 1, \dots, n$, where e_i, e_j are vectors of the canonical basis, so this is not restrictive. On the other hand, nonlinear eigenvalue problems arising in applications are often naturally given in this form with a small k [5, 15]. In particular, the formulation (1) also includes matrix polynomials of degree $k - 1$.

2.1 Backward errors for given eigenpairs

Consider the nonlinear eigenvalue problem $F(\lambda)v = 0$, and assume that we have identified p approximate eigenpairs, for which we have the relations

$$F(\hat{\lambda}_i)\hat{v}_i = \sum_{j=1}^k f_j(\hat{\lambda}_i)F_j\hat{v}_i = r_i.$$

The vectors r_i are the residuals. We provide the formal definition of backward errors for these approximate eigenpairs.

Definition 2.1. Given a nonlinear matrix-valued function $F(\lambda)$, consider p approximate eigenpairs $(\hat{\lambda}_i, \hat{v}_i)$, for $i = 1, \dots, p$. We define the *backward error of the eigenpairs* $(\hat{\lambda}_i, \hat{v}_i)$ as

$$\eta := \min \left\{ \|\delta F_1, \dots, \delta F_k\|_F \mid \sum_{j=1}^k f_j(\hat{\lambda}_i)(F_j + \delta F_j)\hat{v}_i = 0, 1 \leq i \leq p \right\}.$$

Remark 2.2. We note that a trivial solution always exists by taking $\delta F_j = -F_j$, so the minimum is taken on a non-empty set, and that the backward error is always well-defined.

In some frameworks it is important to assign weights to the perturbations. In our setting, this can be achieved considering a weighted matrix-valued function:

$$\alpha_1 f_1(\lambda) F_1 + \dots + \alpha_k f_k(\lambda) F_k,$$

and assigning different values to the weights α_i . For simplicity, we do not include this variant in our results, but a scale of this form is possible.

The backward error η depends on the approximate eigenpairs. We do not explicitly report this dependence to ease the notation, and we assume that they have been fixed throughout this section. We now give an explicit characterization of η .

Theorem 2.3. *Let G, V be the following matrices:*

$$G := \begin{bmatrix} f_1(\hat{\lambda}_1) & \dots & f_k(\hat{\lambda}_1) \\ \vdots & & \vdots \\ f_1(\hat{\lambda}_p) & \dots & f_k(\hat{\lambda}_p) \end{bmatrix}, \quad V := \begin{bmatrix} \hat{v}_1 & \dots & \hat{v}_p \end{bmatrix},$$

and denote by $G \odot^T V^T$ the Khatri-Rao transpose product between G and V^T . Then the backward error η is equal to

$$\eta = \left\| R \left[(G \odot^T V^T)^\dagger \right]^T \right\|_F, \quad R := \sum_{j=1}^k F_j V f_j(\Lambda),$$

where we define the matrix

$$\Lambda := \begin{bmatrix} \hat{\lambda}_1 & & \\ & \ddots & \\ & & \hat{\lambda}_p \end{bmatrix} \in \mathbb{C}^{p \times p}.$$

In particular $\eta \leq \sigma_{\hat{p}}(G \odot^T V^T)^{-1} \|R\|_F$, where $\hat{p}$ is the rank of $G \odot^T V^T$.

Proof. The definition of η involves perturbations δF_j such that the relation $\sum_{j=1}^k f_j(\hat{\lambda}_i)(F_j + \delta F_j)\hat{v}_i = 0$ holds for $i = 1, \dots, p$. This is a linear relation in δF_j , which can be written in matrix form as follows:

$$\left(\underbrace{\begin{bmatrix} f_1(\hat{\lambda}_1)\hat{v}_1^T & \dots & f_k(\hat{\lambda}_1)\hat{v}_1^T \\ \vdots & & \vdots \\ f_1(\hat{\lambda}_p)\hat{v}_p^T & \dots & f_k(\hat{\lambda}_p)\hat{v}_p^T \end{bmatrix}}_{G \odot^T V^T} \otimes I_n \right) \begin{bmatrix} \text{vec}(\delta F_1) \\ \vdots \\ \text{vec}(\delta F_k) \end{bmatrix} = - \begin{bmatrix} \sum_{j=1}^k f_j(\hat{\lambda}_1) F_j \hat{v}_1 \\ \vdots \\ \sum_{j=1}^k f_j(\hat{\lambda}_p) F_j \hat{v}_p \end{bmatrix}. \quad (2)$$

In view of Remark 2.2 the above linear system admits at least a non-trivial solution. The minimum Euclidean norm solution is given by

$$\begin{bmatrix} \text{vec}(\delta F_1) \\ \vdots \\ \text{vec}(\delta F_k) \end{bmatrix} = - \left[(G \odot^T V^T)^\dagger \otimes I_n \right] r, \quad \text{where } r := \begin{bmatrix} \sum_{j=1}^k f_j(\hat{\lambda}_1) F_j \hat{v}_1 \\ \vdots \\ \sum_{j=1}^k f_j(\hat{\lambda}_p) F_j \hat{v}_p \end{bmatrix}. \quad (3)$$

Using the properties of the Kronecker product, we may write

$$\begin{bmatrix} \text{vec}(\delta F_1) \\ \vdots \\ \text{vec}(\delta F_k) \end{bmatrix} = -\text{vec}(R [(G \odot^T V^T)^\dagger]^T).$$

Relation (3) gives the following upper bound on the backward error η :

$$\eta := \left\| \begin{bmatrix} \delta F_1 \\ \vdots \\ \delta F_k \end{bmatrix} \right\|_F \leq \|(G \odot^T V^T)^\dagger \otimes I_n\|_2 \|r\|_2 = \sigma_{\hat{p}}(G \odot^T V^T)^{-1} \|R\|_F,$$

where $\hat{p}$ is the rank of $G \odot^T V^T$. □

The following result shows that the minimal norm perturbations δF_j have a low-rank structure whenever the number of eigenpairs considered is small, that is $p \ll n$.

Lemma 2.4. *The minimal norm backward errors δF_j of Theorem 2.3 can be expressed as $\delta F_j = -RM_j^T$, for appropriate $n \times p$ matrices M_j , where R is the $n \times p$ residual matrix*

$$R = \sum_{j=1}^k F_j V f_j(\Lambda), \quad \Lambda = \begin{bmatrix} \hat{\lambda}_1 & & \\ & \ddots & \\ & & \hat{\lambda}_p \end{bmatrix}.$$

Proof. We denote by $M := (G \odot^T V^T)$, and we partition its pseudoinverse $M^\dagger$ in $n \times 1$ blocks as follows:

$$M^\dagger = \begin{bmatrix} m_{11} & \dots & m_{1p} \\ \vdots & & \vdots \\ m_{k1} & \dots & m_{kp} \end{bmatrix}, \quad m_{ij} \in \mathbb{C}^n.$$

Then, we use the relation $(m \otimes I_n)s = \text{vec}(sm^T)$ and by substituting in the above relation we get

$$\delta F_j = -(r_1 m_{j1}^T + \dots + r_p m_{jp}^T), \quad r_i := \sum_{j=1}^k f_j(\hat{\lambda}_i) F_j \hat{v}_i.$$

Hence, all δF_j are of rank at most p , and can be rewritten as $\delta F_j = -RM_j^T$ for appropriate $n \times p$ matrices M_j . □

Note that the fact that the δF_j are low-rank allows to easily compute their Frobenius and spectral norms (for instance by means of a reduced QR factorization of R and M_j). In addition, the fact that they all share the same left factor R implies that any linear combination of the δF_j still has rank at most p .

2.2 Backward errors for the eigenvalues

If the eigenvectors are not computed or not available, we may give another definition of backward error as follows:

$$\eta := \min_{\hat{v}_i \neq 0} \min \left\{ \left\| [\delta F_1, \dots, \delta F_k] \right\|_F \mid \exists \delta F_j, \sum_{j=1}^k f_j(\hat{\lambda}_i) (F_j + \delta F_j) \hat{v}_i = 0 \right\}.$$

This definition coincides with minimizing Definition 2.1 over all possible choices of eigenvectors $\hat{v}_i$, since we are looking for the closest nonlinear eigenvalue problem with prescribed eigenvalues, and no constrained on the eigenvectors. Using a small abuse of notation, we denote by η the backward error associated with the eigenvalues $\hat{\lambda}_i$, even when the eigenvectors are not computed.

We may provide a version of Theorem 2.3 suited to this scenario.

Theorem 2.5. *For $i = 1, \dots, p$, denote by $\hat{u}_i, \hat{v}_i$ respectively the left and right singular vectors of the matrix $\sum_{j=1}^k f_j(\hat{\lambda}_i) F_j$, associated with the smallest singular value, denoted by $\hat{\sigma}_i$. Let G be the matrix defined in Theorem 2.3 and V be the following matrix:*

$$V := \begin{bmatrix} \hat{v}_1 & \dots & \hat{v}_p \end{bmatrix}.$$

If $G \odot^T V^T$ has rank $\hat{p}$, then we have the following upper and lower bounds for η :

$$\max_{i=1, \dots, p} \left(\frac{\hat{\sigma}_i}{\sqrt{\sum_{j=1}^k |f_j(\hat{\lambda}_i)|^2}} \right) \leq \eta \leq \sigma_{\hat{p}}(G \odot^T V^T)^{-1} \sqrt{p} \max_{i=1, \dots, p} \hat{\sigma}_i.$$

Proof. We start by proving the upper bound for η . We consider the matrix relation

$$\begin{aligned} \begin{bmatrix} \text{vec}(\delta F_1) \\ \vdots \\ \text{vec}(\delta F_k) \end{bmatrix} &= -[(G \odot^T V^T)^\dagger \otimes I_n] \begin{bmatrix} \sum_{j=1}^k f_j(\hat{\lambda}_1) F_j \hat{v}_1 \\ \vdots \\ \sum_{j=1}^k f_j(\hat{\lambda}_p) F_j \hat{v}_p \end{bmatrix} \\ &= -[(G \odot^T V^T)^\dagger \otimes I_n] \begin{bmatrix} \hat{\sigma}_1 \hat{u}_1 \\ \vdots \\ \hat{\sigma}_p \hat{u}_p \end{bmatrix}, \end{aligned}$$

from which we have the following upper bound:

$$\begin{aligned} \eta &\leq \sigma_{\hat{p}}(G \odot^T V^T)^{-1} \left\| \begin{bmatrix} \hat{\sigma}_1 \hat{u}_1 \\ \vdots \\ \hat{\sigma}_p \hat{u}_p \end{bmatrix} \right\|_2 = \sigma_{\hat{p}}(G \odot^T V^T)^{-1} \sqrt{\sum_{i=1}^p \hat{\sigma}_i^2} \\ &\leq \sigma_{\hat{p}}(G \odot^T V^T)^{-1} \sqrt{p} \max_{i=1, \dots, p} \hat{\sigma}_i. \end{aligned}$$

For each $i = 1, \dots, p$, starting from the relation $\sum_{j=1}^k f_j(\hat{\lambda}_i) \delta F_j \hat{v}_i = -\hat{\sigma}_i \hat{u}_i$, we have that:

$$\begin{aligned} \hat{\sigma}_i &= \left\| \sum_{j=1}^k f_j(\hat{\lambda}_i) \delta F_j \hat{v}_i \right\|_F \leq \sum_{j=1}^k |f_j(\hat{\lambda}_i)| \|\delta F_j\|_F \\ &\leq \sqrt{\sum_{j=1}^k |f_j(\hat{\lambda}_i)|^2} \sqrt{\sum_{j=1}^k \|\delta F_j\|_F^2} \\ &\leq \sqrt{\sum_{j=1}^k |f_j(\hat{\lambda}_i)|^2} \eta. \end{aligned}$$

Then maximizing over $i = 1, \dots, p$, we obtain the following lower bound for η :

$$\eta \geq \max_{i=1, \dots, p} \left(\frac{\hat{\sigma}_i}{\sqrt{\sum_{j=1}^k |f_j(\hat{\lambda}_i)|^2}} \right).$$

□

For the case $p = 1$, we obtain an explicit expression for the backward error η , which coincides with the one proposed by Ahmad and Mehrmann (Proposition 2.2, [3]).

Corollary 2.6. *For the case $p = 1$, we have an explicit expression for the backward error*

$$\eta = \frac{\hat{\sigma}_1}{\sqrt{\sum_{j=1}^k |f_j(\hat{\lambda}_1)|^2}},$$

where $\hat{\sigma}_1$ denotes the smallest singular value of the matrix $\sum_{j=1}^k f_j(\hat{\lambda}_1) F_j$.

Proof. Let $\hat{v}_1$ be the right singular vector associated with the singular value $\hat{\sigma}_1$. Then the upper bound for η proposed in Theorem 2.5 may be written as

$$\sigma_{\min}(G \odot^T \hat{v}_1^T)^{-1} \hat{\sigma}_1 = \|G \otimes \hat{v}_1^T\|_2^{-1} \hat{\sigma}_1 = \|G\|_2^{-1} \hat{\sigma}_1 = \frac{\hat{\sigma}_1}{\sqrt{\sum_{j=1}^k |f_j(\hat{\lambda}_1)|^2}}.$$

□

2.3 Explicit upper bounds for the backward errors

The backward errors depend on the norm of the pseudoinverse $G \odot^T V^T$, which is not necessarily easy or cheap to compute. In this section, we provide some upper bounds that can be used in place of computing the norm explicitly.

Lemma 2.7. *Let G be a $p \times k$ matrix and V be a $n \times p$ matrix, with V scaled to have $\|Ve_i\|_2 = 1$ for $i = 1, \dots, p$, where e_i is the i -th vector of the canonical basis. The following bounds for the norm of $\|(G \odot^T V^T)^\dagger\|_2$ hold:*

- If $p \leq kn$, then $\|(G \odot^T V^T)^\dagger\|_2 \leq \sigma_p(G)^{-1} \kappa_2(V)$,
- If $p \leq k$, then $\|(G \odot^T V^T)^\dagger\|_2 \leq \sigma_p(G)^{-1}$,

where $\kappa(V) = \sigma_1(V)/\sigma_p(V)$ is the condition number of V .

Proof. Let us denote by $M := G \odot^T V^T$. We first prove that if $p \leq kn$, then $\|M^\dagger\|_2 \leq \sigma_p(G)^{-1} \kappa_2(V)$. The condition $p \leq kn$ implies that $\|M^\dagger\|_2 = \sigma_p(M)^{-1}$.

Let us denote by J the $p \times p^2$ submatrix of I_{p^2} such that $J(G \otimes V^T) = G \odot^T V^T$; then, we have

$$\sigma_p(M) = \sigma_p(J(G \otimes V^T)) \geq \sigma_{p^2}(G \otimes V^T) = \sigma_p(G) \sigma_p(V),$$

where we have used the J is a unitary projection and therefore for any W it holds $\sigma_{\min}(JW) \geq \sigma_{\min}(W)$, and the properties of the Kronecker product. Since we assumed that $\|Ve_i\|_2 = 1$ which implies $\|V\|_2 \geq 1$, we have $\sigma_p(V) \geq \kappa_2(V)^{-1}$, and we conclude.

To prove the second inequality, consider a QR factorization of G^T , which has the form $Q \begin{bmatrix} R \\ 0 \end{bmatrix} = G^T$ for a $k \times k$ matrix Q and an upper triangular R . We then define X as follows:

$$G^T = Q \begin{bmatrix} R \\ 0 \end{bmatrix}, \quad X = Q \begin{bmatrix} R^{-T} & 0 \\ \sigma_p(G)^{-1} I_{k-p} \end{bmatrix} \implies GX = \begin{bmatrix} I_p & 0 \end{bmatrix}.$$

We now right multiply M by $X \otimes I_n$, which yields a matrix with the following block structure:

$$\hat{M} := M(X \otimes I_n) = \begin{bmatrix} I_p & 0 \end{bmatrix} \odot^T V^T = \begin{bmatrix} v_1^T & & \\ & \ddots & \\ & & v_p^T \end{bmatrix},$$

where we have used the notation $v_i := V e_i$. We have the following relation between singular values of M and $\hat{M}$:

$$\sigma_i(M) \geq \sigma_i(\hat{M}) \sigma_p(X^{-1}) = \sigma_i(\hat{M}) \sigma_1(X)^{-1} = \sigma_i(\hat{M}) \sigma_p(G).$$

We now prove that $\sigma_p(\hat{M}) \geq 1$, which concludes the proof. By the variational characterization of the singular values, we may write

$$\begin{aligned} \sigma_p(\hat{M}) &:= \min_{\|w\|_2=1} \|w^T \hat{M}\|_2 = \|[w_1 v_1^T \quad \dots \quad w_p v_p^T \times \quad \dots \quad \times]\|_2 \\ &\geq \|[w_1 v_1^T \quad \dots \quad w_p v_p^T]\|_2 = \sqrt{\sum_{i=1}^p |w_i|^2 \|v_i\|_2^2} = \sqrt{\sum_{i=1}^p |w_i|^2} = 1, \end{aligned}$$

where in the last steps we have used $\|v_i\|_2 = 1$ and $\|w\|_2 = 1$. □

Remark 2.8. Note that in principle it may happen $\kappa(V) = \infty$ or $\sigma_p(G) = 0$. In both cases, the statement of Lemma 2.7 still holds and yields $\|(G \odot^T V)^\dagger\|_2 \leq \infty$.

3 Structured nonlinear eigenvalue problems

In this section we propose an extension of our analysis that deals with the case when the coefficients F_j have a specific structure that should be preserved in the backward error. For instance, they could be symmetric, Toeplitz, with a given sparsity pattern, or low-rank. Depending on the structure that we consider, we may need to provide a different approach for the computation of the backward error. More specifically, we make the assumption that $F_j \in \mathcal{S}_j \subseteq \mathbb{C}^{n \times n}$, where $\mathcal{S}_j$ is a set of matrices with a particular structure.

We assume that $\mathcal{S}_j$ are at least differentiable manifolds that includes the zero. It is convenient to distinguish two cases:

1. For linear structures, when all the $\mathcal{S}_j$ are linear subspaces, we provide a formula for the structured backward error associated with a set of approximate eigenpairs; we will describe this case in Section 3.1. We will provide some results that hold for symmetric matrices in Section 3.2.
2. For nonlinear structures (such as fixed rank matrices), we propose an approximate upper bound, computed employing a Riemannian optimization-based approach; we will describe it in Section 3.3.

3.1 Structured coefficients in linear subspaces

If the sets $\mathcal{S}_j$ are linear subspaces of $\mathbb{C}^{n \times n}$, then we can write the F_j in an appropriate basis:

$$F_j = \sum_{i=1}^{d_j} \delta_i^j P^{(i,j)}, \quad \mathcal{S}_j = \text{span}(P^{(1,j)}, \dots, P^{(d_j,j)}).$$

From now on, we will assume that the basis given by the $P^{(i,j)}$ are orthogonal with respect to the Frobenius inner product, and normalized to have Frobenius norm equal to 1. This implies that the matrix

$$P^{(j)} = \begin{bmatrix} \text{vec}(P^{(1,j)}), & \dots, & \text{vec}(P^{(d_j,j)}) \end{bmatrix}$$

has orthonormal columns. We denote by P the block diagonal matrix collecting all $P^{(j)}$, defined as follows:

$$P := \begin{bmatrix} P^{(1)} & & \\ & \ddots & \\ & & P^{(k)} \end{bmatrix}. \quad (4)$$

Note that since the $P^{(j)}$ are not square, the above matrix is rectangular as well, and has orthonormal columns. Throughout this section, particular results for the unstructured case can be obtained simply choosing $P = I$.

Definition 3.1. Given $F_j \in \mathcal{S}_j$, for $j = 1, \dots, k$, where $\mathcal{S}_j$ are linear subspaces of $\mathbb{C}^{n \times n}$, consider (V, Λ) , defined as in Theorem 2.3, approximate eigenpairs for the matrix-valued function $F(\lambda) = \sum_{j=1}^k f_j(\lambda) F_j$. The *structured backward error* associated with (V, Λ) is defined as:

$$\eta_{\mathcal{S}} := \min \left\{ \left\| [\delta F_1, \dots, \delta F_k] \right\|_F : \sum_{j=1}^k (F_j + \delta F_j) V f_j(\Lambda) = 0, \delta F_j \in \mathcal{S}_j, \text{ for } j = 1, \dots, k \right\}.$$

We prove the structured analogue of Theorem 2.3.

Theorem 3.2. Let (V, Λ) approximate eigenpairs for the nonlinear eigenvalue problem with structured coefficients $F(\lambda)$, such that

$$R = \sum_{j=1}^k F_j V f_j(\Lambda),$$

and let G be defined as in Theorem 2.3. Then, the structured backward error $\eta_{\mathcal{S}}$ is equal to

$$\eta_{\mathcal{S}} = \left\| \left[((G \odot^T V^T) \otimes I_n) P \right]^{\dagger} r \right\|_2, \quad r := \text{vec}(R),$$

where P is defined as in (4). In particular, we have the upper bound

$$\eta_{\mathcal{S}} \leq \sigma_{\min}(((G \odot^T V^T) \otimes I_n) P)^{-1} \|R\|_F.$$

Proof. Starting with relation (2) provided in Theorem 2.3, we get

$$((G \odot^T V^T) \otimes I_n) \begin{bmatrix} \text{vec}(\delta F_1) \\ \vdots \\ \text{vec}(\delta F_k) \end{bmatrix} = -r.$$

We observe that $\text{vec}(\delta F_j) = P^{(j)} \delta^j$, where $\delta^j = [\delta_1^j \ \dots \ \delta_{d_j}^j]^T$, for $j = 1, \dots, k$ and therefore $\|\delta F_j\|_F = \|\delta^j\|_2$. Then we may write the previous relation as

$$((G \odot^T V^T) \otimes I_n) P \delta = -r, \quad \delta = \begin{bmatrix} \delta^1 \\ \vdots \\ \delta^k \end{bmatrix}.$$

Since $\|\delta\|_2 = \|\begin{bmatrix} \delta F_1 & \dots & \delta F_k \end{bmatrix}\|_F$, we conclude the proof. The upper bound for η_S follows from:

$$\|\delta\|_2 = \left\| \begin{bmatrix} \delta F_1 \\ \vdots \\ \delta F_k \end{bmatrix} \right\|_F \leq \|[(G \odot^T V^T) \otimes I_n] P\|_2 \|r\|_2 = \sigma_{\min}(((G \odot^T V^T) \otimes I_n) P)^{-1} \|R\|_F.$$

□

3.1.1 Invariant pairs

It is possible to provide a generalization of Theorem 3.2 using the notion of invariant pairs. Given a nonlinear matrix-valued function $F(\lambda) = \sum_{j=1}^k F_j f_j(\lambda)$, we say that (V, M) is an invariant pair if the following relation holds:

$$\sum_{j=1}^k F_j V f_j(M) = 0.$$

Note that this implies that $\Lambda(M)$ is a subset of the spectrum of $F(\lambda)$ and that the associated eigenvectors belong to the column span of V . Besides being useful for analyzing (for instance) stable subspaces, this also allows to maintain real arithmetic in case of complex conjugate eigenvalues.

In this setting, denoting by

$$\widehat{G} := [f_1(M)^T \ \dots \ f_k(M)^T], \quad R := \sum_{j=1}^k F_j V f_j(M),$$

and proceeding as in the proof of Theorem 3.2, we have that:

$$\eta_S = \left\| \left[\left(\widehat{G} (I_k \otimes V^T) \right) \otimes I_n \right] P \right\|_2^\dagger \|r\|_2, \quad r := \text{vec}(R),$$

and consequently the upper bound

$$\eta_S \leq \sigma_{\min}(\left(\widehat{G} (I_k \otimes V^T) \right) \otimes I_n) P)^{-1} \|R\|_F.$$

Even though Theorem 3.2 provides an explicit formula for the backward error, the linear system that needs to be solved is much larger than the one in the non-structured case. Hence, it is sometimes convenient to obtain the backward error through the same optimization procedures that we will describe for nonlinear structures in Section 3.3.

3.2 Symmetric backward errors

For particular choices of $\mathcal{S}_j$, we can provide a more detailed analysis. We now focus on the case of real symmetric coefficients $F_j = F_j^T$. For the standard eigenvalue problem, in [20] Tisseur

provides a complete survey on structured backward errors associated with multiple approximate eigenpairs. In particular, a formula is provided for computing the symmetric backward error, using the result on structured matrix problems in [20, Lemma 2.3]. Our result is a generalization to the context of nonlinear eigenvalue problems with symmetric coefficients.

For simplicity, we only discuss the case of eigenvalues and eigenvectors, even though the same analysis can be generalized to invariant pairs with a moderate effort. The problem can be stated as, given a $n \times p$ matrix V and a diagonal matrix Λ , finding real symmetric perturbations $\delta F_j = \delta F_j^T$ such that

$$\sum_{j=1}^k (F_j + \delta F_j) V f_j(\Lambda) = 0. \quad (5)$$

Theorem 3.3. *Let F_j be real symmetric matrices, and (V, Λ) approximate eigenpairs with a diagonal matrix Λ such that $R = \sum_{j=1}^k F_j V f_j(\Lambda)$. Let $V = QT$ be an economy size QR factorization of V , and define*

$$\tilde{T} = \begin{bmatrix} T f_1(\Lambda) \\ \vdots \\ T f_k(\Lambda) \end{bmatrix}, \quad M_S := \begin{bmatrix} \frac{\tilde{T}^T \otimes I_p}{\Pi_{p,p} - I_{p^2}} & & \\ & \ddots & \\ & & \Pi_{p,p} - I_{p^2} \end{bmatrix},$$

where $\Pi_{p,p}$ is the (p, p) commutation matrix (or perfect shuffle, see [24]). Then, there exist symmetric real perturbations δF_j such that (V, Λ) are eigenpairs for $\sum_{j=1}^k (F_j + \delta F_j) f_j(\lambda)$, and

$$\delta F_j = Q \begin{bmatrix} A_{11}^{(j)} & \left(A_{21}^{(j)}\right)^T \\ A_{21}^{(j)} & 0_{(n-p) \times (n-p)} \end{bmatrix} Q^T,$$

where $A_{11}^{(j)}$ and $A_{21}^{(j)}$ solve the equations:

$$\begin{bmatrix} A_{21}^{(1)} & \dots & A_{21}^{(k)} \end{bmatrix} = B_2 \tilde{T}^\dagger, \quad \text{and} \quad M_S \begin{bmatrix} \text{vec}(A_{11}^{(1)}) \\ \vdots \\ \text{vec}(A_{11}^{(k)}) \end{bmatrix} = \begin{bmatrix} \text{vec}(B_1) \\ 0 \\ \vdots \\ 0 \end{bmatrix},$$

where we denote by B_1, B_2 the $p \times p$ and $(n-p) \times p$ blocks of $\begin{bmatrix} B_1 \\ B_2 \end{bmatrix} = -Q^T R$, respectively.

Proof. We note that choosing $\delta F_j = -F_j$ gives a valid solution to the linear system; hence, we know a-priori that the set of all solutions is non-empty, and we look for the minimum norm one. As a preliminary step, we choose a unitary matrix Q such that

$$V f_j(\Lambda) = Q T_j = Q \begin{bmatrix} \tilde{T}_j \\ 0_{(n-p) \times p} \end{bmatrix}, \quad \tilde{T}_j \in \mathbb{C}^{p \times p},$$

with upper triangular matrices $\tilde{T}_j$. Since $f_j(\Lambda)$ is diagonal, such Q can be constructed from a QR factorization of V . We now left multiply (5) by Q^T obtaining

$$\sum_{j=1}^k Q^T \delta F_j Q T_j = - \sum_{j=1}^k Q^T F_j Q T_j =: -Q^T R,$$

where we have used $Q^T V f_j(\Lambda) = T_j$. By partitioning $Q^T \delta F_j Q$ as follows

$$Q^T \delta F_j Q = \begin{bmatrix} A_{11}^{(j)} & (A_{21}^{(j)})^T \\ A_{21}^{(j)} & A_{22}^{(j)} \end{bmatrix}, \quad -Q^T R = \begin{bmatrix} B_1 \\ B_2 \end{bmatrix},$$

we can rewrite the previous equation as

$$\sum_{j=1}^k Q^T \delta F_j Q T_j = \sum_{j=1}^k \begin{bmatrix} A_{11}^{(j)} & (A_{21}^{(j)})^T \\ A_{21}^{(j)} & A_{22}^{(j)} \end{bmatrix} \begin{bmatrix} \tilde{T}_j \\ 0_{(n-p) \times p} \end{bmatrix} = \sum_{j=1}^k \begin{bmatrix} A_{11}^{(j)} \\ A_{21}^{(j)} \end{bmatrix} \tilde{T}_j = \begin{bmatrix} B_1 \\ B_2 \end{bmatrix}.$$

The only condition to have a symmetric solution is to ensure that $A_{11}^{(j)} = (A_{11}^{(j)})^T$ for all $j = 1, \dots, k$, and the above equation decouples in the independent linear systems

$$\sum_{j=1}^k A_{11}^{(j)} \tilde{T}_j = B_1, \quad \sum_{j=1}^k A_{21}^{(j)} \tilde{T}_j = B_2. \quad (6)$$

Note that

$$\|Q^T \delta F_j Q\|_F^2 = \|A_{11}^{(j)}\|_F^2 + 2\|A_{21}^{(j)}\|_F^2 + \|A_{22}^{(j)}\|_F^2.$$

Since we are looking for the minimum norm solution, we can choose $A_{22}^{(j)} = 0$, and $A_{21}^{(j)}$ as the minimum norm solution of the right equation in (6):

$$\begin{bmatrix} A_{21}^{(1)} & \dots & A_{21}^{(k)} \end{bmatrix} = B_2 \tilde{T}^\dagger, \quad \tilde{T} := \begin{bmatrix} \tilde{T}_1 \\ \vdots \\ \tilde{T}_k \end{bmatrix}. \quad (7)$$

To determine $A_{11}^{(j)}$, we write a linear system with the left equation in (6) together with the symmetry condition $A_{11}^{(j)} = (A_{11}^{(j)})^T$. This yields

$$\underbrace{\begin{bmatrix} \tilde{T}^T \otimes I_p \\ \Pi_{p,p} - I_{p^2} \\ \ddots \\ \Pi_{p,p} - I_{p^2} \end{bmatrix}}_{=: M_S} \begin{bmatrix} \text{vec}(A_{11}^{(1)}) \\ \vdots \\ \text{vec}(A_{11}^{(k)}) \end{bmatrix} = \begin{bmatrix} \text{vec}(B_1) \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad (8)$$

where $\Pi_{p,p}$ is the commutation matrix (or perfect shuffle) such that $\Pi_{p,p} \text{vec}(X) = \text{vec}(X^T)$ [24]. We know that the system is solvable, so we can characterize the minimum norm solution by taking the pseudoinverse of the matrix on the left. $\square$

Corollary 3.4. *Under the hypotheses and the notation of Theorem 3.3, we have the following upper bound for the structured backward error associated with the approximate eigenpairs (V, Λ) :*

$$\eta \leq \|R\|_F^2 \left(\|M_S^\dagger\|_F^2 + 2\|\tilde{T}\|_F^2 \right).$$

Proof. Consider the minimum solution δF_j , given by Theorem 3.3. The relation (8) yields to the upper bound $\|A_{11}^{(1)}\|_F^2 + \dots + \|A_{11}^{(k)}\|_F^2 \leq \|B_1\|_F^2 \cdot \|M_S^\dagger\|_F^2$. Combining this and the expression on $A_{21}^{(j)}$ in (7), we obtain

$$\begin{aligned} \sum_{j=1}^k \|\delta F_j\|_F^2 &= \sum_{j=1}^k \|Q^T \delta F_j Q\|_F^2 = \sum_{j=1}^k \left(\|A_{11}^{(j)}\|_F^2 + 2\|A_{21}^{(j)}\|_F^2 \right) \\ &\leq \|R\|_F^2 \left(\|M_S^\dagger\|_F^2 + 2\|\tilde{T}\|_F^2 \right). \end{aligned} \quad \square$$

Remark 3.5. In Corollary 3.4, a bound for the backward error can be computed cheaply, whenever p is small. Indeed, the matrices M_S and $\tilde{T}$ can be computed with $\mathcal{O}(np^2 + p^6 k^3)$. A lower complexity in p may be achieved exploiting the structure of M_S . The dominant term is np^2 as long as $p^4 k^3 < n$, which is realistic in large scale applications where only a few eigenmodes are necessary.

3.3 Nonlinear structures

In this section we describe the more general case of coefficients F_j belonging to a differentiable manifold, which may not be a linear subspace. One of the most relevant examples is taking $F_j \in \mathbb{R}_{r_j}^{n \times n}$, where we denote by $\mathbb{R}_{r_j}^{n \times n}$ the set of real matrices of rank r_j and size $n \times n$.

In this context, we may not write the perturbations δF_j as linear combinations of a set of matrices $P^{(1)}, \dots, P^{(d_j)}$. In this case, we have that each matrix coefficient F_j and each perturbed coefficient $F_j + \delta F_j$ belong to a manifold $\mathcal{S}_j \subseteq \mathbb{R}^{n \times n}$. Given $F_j \in \mathcal{S}_j$, we can rephrase the definition of backward error in the following way:

$$\eta_S = \min \left\{ \|\delta F_1, \dots, \delta F_k\|_F : \sum_{j=1}^k (F_j + \delta F_j) V f_j(\Lambda) = 0, (F_j + \delta F_j) \in \mathcal{S}_j \right\},$$

where (V, Λ) contains approximate eigenpairs, as defined as in Subsection 3.1.

Remark 3.6. Note that the previous definition coincides with the one provided in Subsection 3.1 for linear structures, where we have that $\delta F_j \in \mathcal{S}_j$.

In this setting, we may not provide an explicit formula for the structured backward error, as in Subsection 3.1. Nevertheless, we may numerically approximate an upper bound for the structured backward error, employing Riemannian optimization. Denote by $\tilde{F}_j := F_j + \delta F_j$, for $j = 1, \dots, k$. Consider a $\mu > 0$, we may define the functional

$$\begin{aligned} f : \mathcal{S}_1 \times \dots \times \mathcal{S}_k &\mapsto \mathbb{R} \\ f(\tilde{F}_1, \dots, \tilde{F}_k) &\mapsto \left\| \sum_{j=1}^k \tilde{F}_j V f_j(\Lambda) \right\|_F^2 + \mu \left\| [\tilde{F}_1 - F_1 \quad \dots \quad \tilde{F}_k - F_k] \right\|_F^2. \end{aligned} \tag{9}$$

The parameter μ is needed to force the optimization algorithm to find a minimum norm solution. Therefore, an upper bound for the structured backward error may be obtained minimizing the functional f on the product manifold $\mathcal{S} := \mathcal{S}_1 \times \dots \times \mathcal{S}_k$. This setting allows us to employ Riemannian optimization, minimizing the function f on the product manifold $\mathcal{S}$. We chose to implement this idea relying on `manopt`, a MATLAB package for Riemannian optimization [7], and in particular its implementation of the Riemannian trust region method. To this end, we recall a few results on the product manifolds and perform the computation of the Riemannian gradient and the Riemannian Hessian on $\mathcal{S}$, which we need for the trust-region method on Riemannian manifolds [1].

The product manifold $\mathcal{S}$ can be treated working separately on the manifolds $\mathcal{S}_j$. Indeed, the tangent space of $\mathcal{S}$ can be defined as

$$T_{(\tilde{F}_1, \dots, \tilde{F}_k)}(\mathcal{S}) := T_{\tilde{F}_1}(\mathcal{S}_1) \times \dots \times T_{\tilde{F}_k}(\mathcal{S}_k),$$

and the scalar product that we consider on it is the one inherited from the products on $\mathcal{S}_j$ for $j = 1, \dots, k$, that is:

$$\langle (u_1, \dots, u_k), (w_1, \dots, w_k) \rangle_{(\tilde{F}_1, \dots, \tilde{F}_k)}^{\mathcal{S}} := \langle u_1, w_1 \rangle_{\tilde{F}_1}^{\mathcal{S}_1} + \dots + \langle u_k, w_k \rangle_{\tilde{F}_k}^{\mathcal{S}_k},$$

where $(u_1, \dots, u_k), (w_1, \dots, w_k) \in T_{(\tilde{F}_1, \dots, \tilde{F}_k)}(\mathcal{S})$ and $\langle \cdot, \cdot \rangle^{\mathcal{S}_j}$ is the product associated with $\mathcal{S}_j$. In this setting, we consider only embedded manifolds $\mathcal{S}_j \subseteq \mathbb{R}^{n \times n}$ for $j = 1, \dots, k$, then the product is the real inner product $\langle u, v \rangle := \text{trace}(u^T v)$.

Both the Riemannian gradient and the Riemannian Hessian for the function f can be computed starting from the Euclidean ones. In particular, the Riemannian gradient is obtained computing the orthogonal projection of the Euclidean gradient of f (here we denote by f the smooth extension of the functional (9) to the ambient space) onto the tangent space $T(\mathcal{S})$. The computation of the Riemannian Hessian of f needs both the Euclidean gradient and the Euclidean Hessian for f and it can be obtained through the Weingarten map (see [6, Section 5] for the details).

Even though the projection of both the gradient and the Hessian on the product manifold $\mathcal{S}$ is handled automatically in **manopt**, we will need to implement this carefully to make it efficient. To this end, we first need to derive the Euclidean gradient and the Euclidean Hessian of the functional f .

It is convenient to write the functional as

$$f(\tilde{F}_1, \dots, \tilde{F}_k) = \langle \tilde{F}W, \tilde{F}W \rangle + \mu \langle \tilde{F} - F, \tilde{F} - F \rangle,$$

where

$$\tilde{F} := [\tilde{F}_1 \quad \dots \quad \tilde{F}_k], \quad F := [F_1 \quad \dots \quad F_k] \quad \text{and} \quad W := \begin{bmatrix} Vf_1(\Lambda) \\ \vdots \\ Vf_k(\Lambda) \end{bmatrix}.$$

In order to compute the Euclidean gradient of the functional f , we perform the directional derivative of f :

$$Df(\tilde{F})[\tilde{E}] = \left. \frac{d}{dt} f(\tilde{F}_1 + t\tilde{E}_1, \dots, \tilde{F}_k + t\tilde{E}_k) \right|_{t=0},$$

in the direction $\tilde{E} := [\tilde{E}_1 \quad \dots \quad \tilde{E}_k]$. In this case, we may write

$$\begin{aligned} \left. \frac{d}{dt} f(\tilde{F}_1 + t\tilde{E}_1, \dots, \tilde{F}_k + t\tilde{E}_k) \right|_{t=0} &= \sum_{j=1}^k 2 \langle \tilde{E}_j W_j, \tilde{F}W \rangle + 2\mu \langle \tilde{E}_j, \tilde{F}_j - F_j \rangle \\ &= \sum_{j=1}^k \langle \tilde{E}_j, 2\tilde{F}W W_j^T + 2\mu(\tilde{F}_j - F_j) \rangle, \end{aligned}$$

where W_j is the j -th block row of W and we used the circulant property of the trace in the last step. Then using that the Euclidean gradient is the unique vector such that

$$\forall \tilde{F}, \tilde{E} \in \mathbb{R}^{n \times n} \times \dots \times \mathbb{R}^{n \times n} \quad Df(\tilde{F})[\tilde{E}] = \langle \tilde{E}, \text{grad} f(\tilde{F}) \rangle,$$

we conclude that $\text{grad} f(\tilde{F}) = 2\tilde{F}W W^T + 2\mu(\tilde{F} - F)$. In the implementation of the Euclidean gradient in **manopt**, it is useful to split the contributions for each term of the product manifold. Then we may consider $\text{grad}_{\tilde{F}_j} f(\tilde{F}) = 2\tilde{F}W W_j^T + 2\mu(\tilde{F}_j - F_j)$.

The Euclidean Hessian of the function $f : \mathbb{R}^{n \times n} \times \dots \times \mathbb{R}^{n \times n} \mapsto \mathbb{R}$ at the point $(\tilde{F}_1, \dots, \tilde{F}_k)$ is defined as the directional derivative of the Euclidean gradient $\text{grad} f$

$$\text{Hess} f(\tilde{F})[\tilde{E}] = D \text{grad} f(\tilde{F})[\tilde{E}] = \lim_{t \rightarrow 0} \frac{\text{grad} f(\tilde{F} + t\tilde{E}) - \text{grad} f(\tilde{F})}{t},$$

where $\tilde{E} \in \mathbb{R}^{n \times n} \times \dots \times \mathbb{R}^{n \times n}$. Note that here we denote by f both the functional (9) and its smooth extension to the ambient space. Inside the Riemannian Trust Region scheme, we only need to evaluate the Euclidean Hessian along a specified direction $\tilde{E} = (\tilde{E}_1, \dots, \tilde{E}_k)$, which is given by the directional derivative

$$\left. \frac{d}{dt} \text{grad} f(\tilde{F}_1 + t\tilde{E}_1, \dots, \tilde{F}_k + t\tilde{E}_k) \right|_{t=0} = 2\tilde{E}WW^T + 2\mu\tilde{E}.$$

Moreover, observing that

$$\text{D grad} f(\tilde{F})[\tilde{E}] := \left(\text{D grad}_{\tilde{F}_1} f(\tilde{F})[\tilde{E}], \dots, \text{D grad}_{\tilde{F}_k} f(\tilde{F})[\tilde{E}] \right),$$

we may split the contributions of the different terms of the product manifold for the implementation in `manopt`, obtaining that

$$\text{D grad}_{\tilde{F}_j} f(\tilde{F})[\tilde{E}] = 2\tilde{E}WW_j^T + 2\mu\tilde{E}_j.$$

Note that both the (Euclidean) gradient and the (Euclidean) Hessian have a first term which is low-rank. Indeed both $2\tilde{F}WW^T$ and $2\tilde{E}WW^T$ are expressed in a low-rank format, therefore for several choices of manifolds we may compute their projection directly in an efficient way.

In Subsection 4.3, we test this approach for a selected number of structures. In particular, we consider the case of sparse matrices, multiples of the identity and fixed rank matrices. For these structures, once we have computed the matrix $\tilde{F}W$, we handle the projection of the term $\tilde{F}WW_j^T$ as follows:

1. $\mathcal{S}_j$ is the set of sparse matrices in $\mathbb{R}^{n \times n}$: let $\mathcal{J} \subseteq \{1, \dots, n\}^2$ the set of indices corresponding to the nonzero entries of $\mathcal{S}_j$, then the matrix-matrix multiplication between $\tilde{F}W$ and W_j^T as

$$(\tilde{F}WW_j^T)_{(a,b)} = \begin{cases} \sum_{c=1}^p (\tilde{F}W)_{(a,c)} (W_j)_{(b,c)} & (a,b) \in \mathcal{J} \\ 0 & (a,b) \notin \mathcal{J} \end{cases}.$$

The complexity for this product is $\mathcal{O}(p|\mathcal{J}|)$.

2. $\mathcal{S}_j$ is the set of matrices multiple of the identity in $\mathbb{R}^{n \times n}$: we can perform the projection of the matrix $\tilde{F}WW_j^T$ onto the tangent space of this manifold simply computing the $\frac{1}{n}\text{trace}(\tilde{F}WW_j^T)$, for which the computational cost is $\mathcal{O}(np^2 + n)$;
3. $\mathcal{S}_j$ is the set of matrices of fixed rank r_j in $\mathbb{R}^{n \times n}$: a rank p matrix is represented as USV^T by storing a structure with three fields $U, V \in \mathbb{R}^{n \times p}$, $S \in \mathbb{R}^{p \times p}$, where U, V are orthonormal and the matrix S is any diagonal or full-rank matrix. The term $\tilde{F}WW_j^T$ can be represented in this way by the matrices $\tilde{F}W$, W_j and I_p , respectively. The latter can be projected on the tangent space of $\mathbb{R}_{r_j}^{n \times n}$ by `manopt` using an economy-size SVD, which requires $\mathcal{O}(nr_j(p + r_j))$ flops.

The same procedures can be repeated for handling the projection of the matrix $\tilde{E}WW_j^T$, for these three manifolds.

In the numerical implementation of the method, we successively solve minimization problems in the form (9), for different choices of the parameter μ . This approach, also known in optimization theory as penalization method, consists in solving the problem for smaller and smaller choices of the parameter μ , using the solution of one step as initial point for the following one. An overview on these call of solvers for constrained optimization is contained in [4, Section 4], while results on their generalization to Riemannian manifolds can be found in [18].

We describe the approach in Algorithm 1, where each minimization problem needs to be solved using the Riemannian based-method proposed in this Section.

Algorithm 1 Riemannian optimization-based algorithm

Input: Matrices $F := [F_1 \cdots F_k]$, manifold $\mathcal{S}$, functions f_i , (V, Λ) approximate eigenpairs, desired accuracy ϵ

Output: $\eta_{\mathcal{S}}$ upper bound for the structured backward error, decrease factor ρ

```
1: Begin
2:   Set  $\mu = 1$ 
3:   Set starting point  $\tilde{F} = F$ 
4:   while  $\sqrt{\mu} > \epsilon$  do
5:      $\tilde{F} \leftarrow \arg \min_{F \in \mathcal{S}} f$ , in (9)
6:      $\mu \leftarrow \rho\mu$ 
7:   end while
8:    $\eta_{\mathcal{S}} \leftarrow \|\tilde{F} - F\|_F$ 
9: End
```

Remark 3.7. In Line 6 of Algorithm 1, we suggest decreasing the parameter μ by a constant factor ρ at each step. Our choice for the factor ρ and additional implementation details are available at the Github repository <https://github.com/miryamgnazzo/backward-error-nonlinear>.

It is common to measure the backward errors using the Euclidean distance, which means computing the norm $\left\| \begin{bmatrix} \tilde{F}_1 - F_1 & \cdots & \tilde{F}_k - F_k \end{bmatrix} \right\|_F$, as we propose in our analysis. However, in some situations it is possible to measure the distance between points on the manifold, using the Riemannian distance. Following the idea in [14, Theorem 3.1], we could rephrase our upper bounds, working directly on the distance between points on the manifold. Since this result employ lower bounds on sectional curvatures, it strictly depends on the geometry of the manifold we use. Indeed, while for flat manifolds the sectional curvature is zero, for different manifolds it can be derived from principal curvatures. In particular, for the case of manifolds of matrices of fixed rank, it can be obtained using [11, Theorem 24].

4 Numerical experiments

This section is devoted to assessing the quality of the theoretical bounds, and to check the effectiveness of the Riemannian optimization scheme in computing the backward errors. We also include tests for symmetric nonlinear eigenvalue problems as described in Section 3.2. For the case of nonlinear structures, our implementation of pseudocode 1 in MATLAB is freely available at <https://github.com/miryamgnazzo/backward-error-nonlinear>, together with the codes for the bounds in Section 2 and Subsections 3.1, 3.2.

Throughout this section, all nonlinear problems for which we need a few eigenvalues to test have been solved with the Newton method initialized with different starting points. The experiments were run using MATLAB 2022b on Intel Core i7-1070H.

4.1 Unstructured tests

4.1.1 The Hader problem

We consider the nonlinear eigenvalue problem in the form:

$$F(\lambda)v = [(e^\lambda - 1)A_2 + \lambda^2 A_1 - \alpha A_0]v = 0, \quad (10)$$

where the coefficient matrices $A_i \in \mathbb{R}^{8 \times 8}$ are symmetric and $\alpha = 100$. This example is known as the Hader problem [13] and it is part of the collection of nonlinear eigenvalue problems in the MATLAB package `nlvep` [5]. We consider a set of $p = 3$ approximate eigenpairs of (10) and randomly generate a set of 1000 perturbation matrices δA_j for $j = 0, 1, 2$. Then we may compute the backward errors using the formula in Theorem 2.3 and test the upper bounds for the unstructured backward error provided in Theorem 2.3 and Lemma 2.7.




Figure 1: Comparison among the upper bounds for the unstructured backward error in Theorem 2.3 and Lemma 2.7, applied to the Hader problem (10).

4.1.2 The beam problem

We consider the delay eigenvalue problem obtained through the finite difference discretization of a one-dimensional beam with delayed stabilizing feedback, as described in [23]:

$$D(\lambda) = -\lambda I_n + A_0 + e^{-\lambda} A_1, \quad n = 1000, \quad (11)$$

where we have

$$A_0 = \begin{bmatrix} A & -w^T \\ -nw & n \end{bmatrix}, \quad A = \text{tridiag}(1, -2, 1) \in \mathbb{R}^{(n-1) \times (n-1)}, \quad w = [0 \quad \dots \quad 1] \in \mathbb{R}^{1 \times (n-1)}$$

and $A_1 = e_n e_n^T$ with e_n the n -th vector of the canonical basis in $\mathbb{R}^n$. The coefficient matrices for this problem can be found in the example gallery presented in the NEP-PACK collection [15]. In Figure 2, we provide a comparison among the upper bounds for the unstructured backward error provided in Theorem 2.3 and Lemma 2.7. It is worth noting that the first two upper bounds for the case $p = 3$ do not coincide, however this is not perceptible on the figure. Observe that for the case $p = 10$, we may not use the first upper bound provided in Lemma 2.7.




Figure 2: Comparison of the upper bounds for the unstructured backward error for the beam problem in (11). On the left: we consider $p = 3$ approximated eigenpairs. On the right: we consider $p = 10$ approximated eigenpairs.

4.1.3 Test on randomly generated problems

We now generate a set of random problems of the form

$$F(\lambda) = A_0 + \lambda A_1 + \lambda^2 I + e^{-\lambda} E_1 + e^{-2\lambda} E_2 \quad (12)$$

where A_0, A_1, E_1, E_2 are randomly symmetric generated matrices. We fix a set of matrices $A_0, A_1, E_1, E_2 \in \mathbb{R}^{128 \times 128}$ of random matrices and a set of approximate eigenpairs $(\hat{\lambda}_i, \hat{v}_i)$ for $i = 1, \dots, p$. Then we generate 1000 random perturbed matrix-valued function in the form:

$$\tilde{F}(\lambda) = \tilde{A}_0 + \lambda \tilde{A}_1 + \lambda^2 \tilde{I} + e^{-\lambda} \tilde{E}_1 + e^{-2\lambda} \tilde{E}_2,$$

and compute the backward error for the approximate eigenpairs $(\hat{\lambda}_i, \hat{v}_i)$ of the nonlinear eigenvalue problem associated with $\tilde{F}(\lambda)$. Then we test the error bounds for the unstructured backward error associated in Theorem 2.3, and compare it with the explicit bounds obtained in Lemma 2.7. In Figure 3, the plot on the left provides the comparison for $p = 3$, while the plot on the right a comparison of the bounds for $p = 10$. Observe that if the number of considered approximate eigenpairs p is strictly larger than the number of coefficients in the matrix-valued function, in Lemma 2.7 the second bound does not hold.

4.2 Structured case: linear subspaces

4.2.1 Randomly generated and sparse matrices

We consider again the matrix-valued function in (12), with randomly generated matrices of size 64×64 , and we impose a sparsity pattern on the coefficients A_0, A_1, E_1, E_2 , where we allow the sparsity patterns to be different from each others. We generate 1000 random sets of coefficients for the matrix-valued function, where we preserve the sparsity pattern on the coefficients. We




Figure 3: Comparison among the upper bounds for the unstructured backward error for the problem (12). On the left: we consider a set of $p = 3$ approximate eigenpairs. On the right: we consider a set of $p = 10$ approximate eigenpairs.

compute the structured backward error associated with a set of $p = 3$ approximate eigenpairs for this set of randomly generated family of matrices, using the result in Theorem 3.2. In Figure 4, we provide test the upper bound for the structured backward error imposing sparsity patterns on the coefficients, against the exact formula for the structured backward error, as provided in Theorem 3.2. We report for completeness the upper bound for unstructured backward error provided by 2.3 (which does not hold in this case).

4.2.2 Randomly generated symmetric matrices

Consider again the nonlinear eigenvalue problem associated with (12), with randomly generated coefficients such that $A_i = A_i^T$ for $i = 0, 1$ and $E_j = E_j^T$ for $j = 1, 2$. As in the previous case, we run 1000 tests, for a set of $p = 3$ approximated eigenpairs. In Figure 5, we consider an example of size $n = 64$ and compute the structured backward error imposing the symmetry on the coefficient matrices, provided in Theorem 3.3. We provide a comparison among the upper bound for general linear structures in Theorem 3.2 and the one specialized for symmetry structures in Corollary 3.4.

Then we consider two randomly generated and symmetric problems as in (12), where the dimension of the coefficients is $n = 128$ and $n = 2048$. In Figure 6, we test the upper bound in Corollary 3.4 against the structured backward error obtained using Theorem 3.3, comparing it with the one for unstructured backward error.




Figure 4: Test the upper bound for the structured backward error in Theorem 3.2, for the case of randomly generated sparse matrices in Subsection 4.2.1. The bound for unstructured case, which does not hold, is reported for completeness.

4.3 Riemannian optimization

4.3.1 Quadratic polynomial eigenvalue problem

We consider the nonlinear matrix-valued function:

$$F(\lambda) = A_0 + \lambda A_1 + \lambda^2 A_2 \in \mathbb{R}^{n \times n}, \quad n = 10000, \quad (13)$$

where the matrix $A_0 = \text{tridiag}(1, -2, 1)$, $A_1 = -UU^T$ is a low-rank matrix with a randomly generated matrix $U \in \mathbb{R}^{10000 \times 2}$ and A_2 is the identity matrix. We consider an approximation of two eigenpairs $(\hat{\lambda}_i, \hat{v}_i)$ for $i = 1, 2$ and perturb the matrix coefficients keeping the same structures:

$$\tilde{F}(\lambda) = (A_0 + \tilde{A}_0) + \lambda \tilde{A}_1 + \lambda^2 (A_2 + \tilde{A}_2),$$

where $\tilde{A}_0$ is a randomly generated tridiagonal matrix, $\tilde{A}_1 = -(U + \tilde{U})(U + \tilde{U})^T$ with $\tilde{U} \in \mathbb{R}^{n \times 2}$ randomly generated and $\tilde{A}_2$ a multiple of the identity. The norm of the perturbation $\|\tilde{F} - F\|_F$ is in the order of 1.992017.

In order to apply the method proposed in Subsection 3.3, we consider the following product manifold:

$$\mathcal{S} := \mathcal{S}_1 \times \mathbb{R}_2^{n \times n} \times \mathcal{S}_2,$$

where $\mathcal{S}_1$ is the manifold of sparse matrices with the same sparsity patterns of A_0 , $\mathbb{R}_2^{n \times n}$ is the manifold of rank 2 real matrices of size $n \times n$ and $\mathcal{S}_2$ is the manifold of the matrices that are multiples of the identity. The implementation of the method requires the use of the `manopt` package for MATLAB, version 7.1. Observe that the manifold $\mathcal{S}_2$ is not available in `manopt`, then we used our implementation of this manifold.




Figure 5: Comparison between the bounds for structured backward error in Theorem 3.2 and Corollary 3.4, applied to problem (12) with symmetric coefficients. For completeness, we report the unstructured bound.




Figure 6: Test the structured bound in Corollary 3.4, for randomly generated symmetric problems in (12). On the left: size $n = 128$. On the right: size $n = 2048$. The bound for unstructured case, which does not hold, is reported for completeness.

The running time for the computation of the structured backward error associated with two

approximate eigenpairs of $F(\lambda)$ is 199.1613 seconds. We obtain an upper bound for the backward error equal to 3.521105×10^{-2} and a norm of the residual equal to 4.015467×10^{-9} .

In practice, this experiment can be repeated as coded in Listing 1, where $\mathbf{f}$ is a function for the evaluation of $[1, \lambda, \lambda^2]$ and $(\mathbf{V}, \mathbf{L})$ are approximate eigenpairs. The command `be_riemannian` calls the manifolds that we need for the optimization procedure, where 'identity' refers to our implementation of the manifold $\mathcal{S}_2$. The MATLAB functions can be found in the github repository <https://github.com/miryamgnazzo/backward-error-nonlinear>.

Listing 1: Code for the experiment in (13)

```
F = { A0, A1, A2 }; %cell array of coefficient matrices

D = be_riemannian(F, @f, ...
    { 'sparse', 'low-rank', 'identity' }, V, L);
nrm = be_norm(D); %Computed backward error;
```

4.3.2 The beam problem with prescribed sparsity pattern

Consider again the beam problem stated in Subsection 4.1.2. We consider different dimensions n for the matrix-valued function in (11) and we apply the Riemannian optimization-based approach in Section 3.3, preserving the sparsity structures of the coefficients, that leads to the product manifold:

$$\mathcal{S} := \mathcal{S}_0 \times \mathcal{S}_1 \times \mathcal{S}_2,$$

where $\mathcal{S}_0$ is the manifold of matrices that are multiples of the identity, $\mathcal{S}_1$ is the manifold of tridiagonal matrices and $\mathcal{S}_2$ the one of multiples of the matrix $e_n e_n^T$. Observe that the involved structures are linear, nevertheless we compute an approximation of the structured backward error, in order to provide a few examples on matrices of large size.

We compute $\hat{V} \in \mathbb{R}^{n \times 3}$ and $\hat{\Lambda} \in \mathbb{R}^{3 \times 3}$, approximations of $p = 3$ eigenvectors and eigenvalues of $D(\lambda)$ in (11), respectively, then perturb it by

$$\Delta D(\lambda) = -\Delta_0 \lambda + \Delta_1 + \Delta_2 e^{-\lambda},$$

where $F_i + \Delta_i \in \mathcal{S}_i$ for $i = 0, 1, 2$, since $F_0 = I_n$, $F_1 = A_0$ and $F_2 = A_1$. The algorithm in Subsection 3.3 provides final matrices $\Delta F_1, \Delta F_2, \Delta F_3$, which we use to define the (approximated) structured backward error $\eta_S = \left\| \begin{bmatrix} \Delta F_0 & \Delta F_1 & \Delta F_2 \end{bmatrix} \right\|_F$. We test the accuracy of our solution computing the norm of the residual

$$R := -(I_n + \Delta F_0) \hat{V} \hat{\Lambda} + (A_0 + \Delta F_1) \hat{V} + (A_1 + \Delta F_2) \hat{V} \exp(-\hat{\Lambda}).$$

In Table 1, we collect the results obtained considering different sizes n . In particular, we provide a comparison among the elapsed time (expressed in seconds), the (approximated) structured backward error η_S , the norm of the residual R and the Frobenius norm of the starting perturbation matrices.

Conclusions

We propose a backward error analysis for nonlinear eigenvalue problems given in split form. We presented a novel formula for the computation of the backward errors for a given set of eigenpairs or eigenvalues, and explicitly, computable, and inexpensive upper bounds for them.

n	time (in seconds)	η_S	$\ R\ _F$	$\ \Delta D\ _F$
10^3	52.9975	5.564350×10^{-3}	2.351591×10^{-9}	5.625286×10^{-3}
2×10^3	62.2178	7.798143×10^{-3}	9.967606×10^{-10}	7.811920×10^{-3}
5×10^3	94.1982	1.383995×10^{-2}	3.560193×10^{-9}	1.421781×10^{-2}
10^4	145.4670	1.814495×10^{-2}	4.786304×10^{-9}	1.859965×10^{-2}
2×10^4	249.8171	2.584339×10^{-2}	2.480937×10^{-9}	2.616653×10^{-2}
5×10^4	945.6344	3.759764×10^{-2}	6.855576×10^{-9}	3.877037×10^{-2}
10^5	1.6083×10^3	5.295991×10^{-2}	1.012763×10^{-8}	5.703946×10^{-2}

Table 1: Results for the beam problem (11), with different sizes n .

These bounds have been verified to be tight and descriptive on a set of examples arising from standard benchmark collections.

We discussed in detail how to impose different structures on the backward errors. For the case of coefficients living in a linear subspace, we have extended the previous analysis, and provided computable bounds. The bounds are in particular still relatively inexpensive for the relevant case of symmetric coefficients.

For more general structures, where coefficients are in a differentiable manifold, we have provided an effective algorithm for the computation of the backward error, based on a Riemannian optimization technique. This allows to compute backward errors for problems with low-rank coefficients, but also for the ones where the constraint is linear, such as prescribed sparsity pattern, or symmetries, and any combination of these. We have verified the effectiveness and the scalability of this approach, which is able to give explicit bounds for large-scale structured nonlinear eigenvalue problems.

Acknowledgments

Miryam Gnazzo and Leonardo Robol are members of the INdAM Research group GNCS (Gruppo Nazionale di Calcolo Scientifico). The work of Leonardo Robol has been supported by the National Research Center in High Performance Computing, Big Data and Quantum Computing (CN1 – Spoke 6), by the MIUR Excellence Department Project awarded to the Department of Mathematics, University of Pisa, CUP I57G22000700001, and by the PRIN 2022 Project “Low-rank Structures and Numerical Methods in Matrix and Tensor Computations and their Application”.

Declarations

Conflict of interest The authors have no competing interests to declare that are relevant to the content of this article.

References

- [1] P. A. Absil, C. G. Baker, and K. A. Gallivan. Trust-region methods on riemannian manifolds. *Found. Comput. Math.*, 7(3):303–330, jul 2007.
- [2] Bibhas Adhikari, Rafikul Alam, and Daniel Kressner. Structured eigenvalue condition numbers and linearizations for matrix polynomials. *Linear Algebra and its Applications*, 435(9):2193–2221, 2011.
- [3] Sk. Safique Ahmad and Volker Mehrmann. Backward errors and pseudospectra for structured nonlinear eigenvalue problems. *Operators and Matrices*, 10:539–556, 09 2016.
- [4] D.P. Bertsekas. *Nonlinear Programming*. Athena Scientific, Belmont, 1999.
- [5] Timo Betcke, Nicholas J. Higham, Volker Mehrmann, Christian Schröder, and Françoise Tisseur. Nlevp: A collection of nonlinear eigenvalue problems. *ACM Trans. Math. Softw.*, 39(2), feb 2013.
- [6] Nicolas Boumal. *An introduction to optimization on smooth manifolds*. Cambridge University Press, 2023.
- [7] Nicolas Boumal, Bamdev Mishra, P.-A. Absil, and Rodolphe Sepulchre. Manopt, a Matlab toolbox for optimization on manifolds. *Journal of Machine Learning Research*, 15(42):1455–1459, 2014.
- [8] Robert M. Corless. Pseudospectra of exponential matrix polynomials. *Theoretical Computer Science*, 479:70–80, 2013.
- [9] Froilan Dopico, Piers Lawrence, Javier Pérez, and Paul Van Dooren. Block Kronecker linearizations of matrix polynomials and their backward errors. *Numerische Mathematik*, 140, 10 2018.
- [10] Froilan Dopico and Kenet Pomés Portal. Structured eigenvalue condition numbers for parameterized quasiseparable matrices. *Numerische Mathematik*, 134, 11 2016.
- [11] Florian Feppon and Pierre F. J. Lermusiaux. A geometric approach to dynamical model order reduction. *SIAM Journal on Matrix Analysis and Applications*, 39(1):510–538, 2018.
- [12] Stefan Güttel and Françoise Tisseur. The nonlinear eigenvalue problem. *Acta Numerica*, 26:1–94, 2017.
- [13] K. P. Hadeler. Mehrparametrisierte und nichtlineare eigenwertaufgaben. *Archive for Rational Mechanics and Analysis*, 27(4):306–328, 1967.
- [14] Simon Jacobsson, Raf Vandebril, Joeri Van der Veken, and Nick Vannieuwenhoven. Approximating maps into manifolds with lower curvature bounds. *ArXiv*, abs/2403.16785, 2024.
- [15] Elias Jarlebring, Max Bannedich, Giampaolo Mele, Emil Ringh, and Parikshit Upadhyaya. Nep-pack: A julia package for nonlinear eigenproblems - v0.2, 2018.
- [16] Michael Karow, Daniel Kressner, and Emre Mengi. Nonlinear eigenvalue problems with specified eigenvalues. *SIAM Journal on Matrix Analysis and Applications*, 35(3):819–834, 2014.

- [17] Pieter Lietaert, Karl Meerbergen, Javier Pérez, and Bart Vandereycken. Automatic rational approximation and linearization of nonlinear eigenvalue problems. *IMA Journal of Numerical Analysis*, 42(2):1087–1115, 02 2021.
- [18] Changshuo Liu and Nicolas Boumal. Simple algorithms for optimization on Riemannian manifolds with constraints. *Appl. Math. Optim.*, 82(3):949–981, dec 2020.
- [19] Anshul Prajapati and Punit Sharma. Structured eigenvalue backward errors for rational matrix functions with symmetry structures. *BIT Numerical Mathematics*, 64, 02 2024.
- [20] Françoise Tisseur. A chart of backward errors for singly and doubly structured eigenvalue problems. *SIAM Journal on Matrix Analysis and Applications*, 24(3):877–897, 2003.
- [21] Françoise Tisseur. Backward error and condition of polynomial eigenvalue problems. *Linear Algebra and its Applications*, 309(1):339–361, 2000.
- [22] Françoise Tisseur and Karl Meerbergen. The quadratic eigenvalue problem. *SIAM Review*, 43(2):235–286, 2001.
- [23] Roel Van Beeumen, Elias Jarlebring, and Wim Michiels. A rank-exploiting infinite Arnoldi algorithm for nonlinear eigenvalue problems. *Numerical Linear Algebra with Applications*, 23(4):607–628, 2016.
- [24] Charles F Van Loan. The ubiquitous Kronecker product. *Journal of computational and applied mathematics*, 123(1-2):85–100, 2000.