



STRUCTURED LEVEL-2 CONDITION NUMBERS OF MATRIX FUNCTIONS*

BAHAR ARSLAN[†], SAMUEL D. RELTON[‡], AND MARCEL SCHWEITZER[§]

Abstract. Matrix functions play an increasingly important role in many areas of scientific computing and engineering disciplines. In such real-world applications, algorithms working in floating-point arithmetic are used for computing matrix functions and additionally input data might be unreliable, e.g., due to measurement errors. Therefore, it is crucial to understand the sensitivity of matrix functions to perturbations, which is measured by condition numbers. However, the condition number itself might not be computed exactly as well due to round-off and errors in the input. The sensitivity of the condition number is measured by the so-called *level-2 condition number*. For the usual (level-1) condition number, it is well known that *structured* condition numbers (i.e., where only perturbations are taken into account that preserve the structure of the input matrix) might be much smaller than unstructured ones, which, e.g., suggests that structure-preserving algorithms for matrix functions might yield much more accurate results than general-purpose algorithms. In this work, we present a novel upper bound on the structured level-2 condition number, focusing on perturbation matrices within an automorphism group, a Lie or Jordan algebra, or the space of quasi-triangular matrices. In numerical experiments, we then compare the unstructured level-2 condition number with the structured one for some specific matrix functions such as the matrix logarithm, matrix square root, and matrix exponential.

Key words. Level-2 condition number, Matrix function, Automorphism group, Lie algebra, Jordan algebra, Quasi-triangular matrices.

AMS subject classifications. 65F35, 15A16, 65F60.

1. Introduction. Matrix functions $f : \mathbb{K}^{n \times n} \rightarrow \mathbb{K}^{n \times n}$, where $\mathbb{K} = \mathbb{C}$ or \mathbb{R} , such as the matrix exponential, logarithm, and square root are used in a wide variety of applications including the solution of differential equations by exponential integrators, network analysis, and machine learning [4, 6, 8, 16, 20, 23, 28]. The condition number measures the sensitivity of a matrix function to perturbations in the input due to rounding or measurement error.

The (level-1) absolute condition number is defined for a matrix function $f : \mathbb{K}^{n \times n} \rightarrow \mathbb{K}^{n \times n}$ as

$$(1.1) \quad \text{cond}^{[1]}(f, A) = \lim_{\epsilon \rightarrow 0} \sup_{\|E\| \leq \epsilon} \frac{\|f(A + E) - f(A)\|}{\epsilon}.$$

It can be computed via the Fréchet derivative $L_f^{(1)}(A, \cdot)$ of the matrix function, a linear function in the perturbation matrix E such that $L_f^{(1)}(A, E)$ satisfies [19, Section 3]

$$\|f(A + E) - f(A) - L_f^{(1)}(A, E)\| = o(\|E\|),$$

*Received by the editors on April 27, 2023. Accepted for publication on December 23, 2023. Handling Editor: Vanni Noferini. Corresponding Author: Marcel Schweitzer.

[†]Mathematics Department, Faculty of Engineering and Natural Sciences, Bursa Technical University, Yıldırım/Bursa, Turkey (bahar.arslan@btu.edu.tr).

[‡]Faculty of Medicine and Health, University of Leeds, Leeds, UK (s.d.relton@leeds.ac.uk).

[§]School of Mathematics and Natural Sciences, Bergische Universität Wuppertal, 42097 Wuppertal, Germany (marcel@uni-wuppertal.de).

as we outline at the beginning of Section 2.1. When the input matrix A has some structure described by a Lie or Jordan algebra (e.g., symmetric or skew-Hermitian), and we consider only perturbations that retain this structure, the resulting *structured level-1 condition number* can be significantly smaller. This has been investigated in [31] for linear systems, matrix inversion, and distance to singularity. The effect of structured perturbations of matrix functions in Lie and the Jordan algebras is also studied by Davies [10]. The structured level-1 condition number using differentials between smooth manifolds is also analyzed for automorphism groups in [7].

Unfortunately, the condition number itself might not be computed precisely: previous work on the condition number of the condition number (the so-called *level-2 condition number*) [22] has shown that it can be highly sensitive in some cases. Similar findings are also known for the level-2 condition number for solving linear systems [13, 17].

The aim of this work is to investigate the structured level-2 condition number of matrix functions by enforcing structure on the perturbation matrices and it is organized as follows. In Section 2, we summarize existing work on the unstructured level-2 condition number of matrix functions and the structured level-1 condition number. We then build upon these foundations in Section 3 and derive upper bounds for structured condition numbers of matrix functions of different structures. We first investigate the case of matrices coming from Lie and Jordan algebras or automorphism groups and then focus on the important case of quasi-triangular matrices. Additionally, for a few specific choices of function and matrix class, we derive exact formulas for the structured level-2 condition number. Numerical experiments in Section 4 provide a thorough comparison of the structured and unstructured condition numbers for the matrix exponential, logarithm, and square root over a variety of structures. We summarize our findings in Section 5 and give some ideas for future work in this area.

2. Previous work on matrix function conditioning. In this section, we summarize the current theory on matrix function conditioning, which we build upon to derive corresponding results for the structured level-2 condition number in later sections. In particular, we focus here on the structured level-1 condition number and the unstructured level-2 condition number. We reproduce some of the derivations in quite a bit of detail, as the concepts and notations introduced there are essential for understanding the derivations of our new results in Section 3.

2.1. The unstructured level-2 condition number case. In this section, we summarize previous results for the unstructured level-2 condition number, largely based upon [19] and [22]. To begin, the level-1 condition number is the norm of the first-order Fréchet derivative operator $L_f^{(1)}(A, \cdot)$, i.e.,

$$\text{cond}^{[1]}(f, A) = \max_{E \neq 0} \frac{\|L_f^{(1)}(A, E)\|}{\|E\|} = \max_{\|E\|=1} \|L_f^{(1)}(A, E)\|,$$

where $\|\cdot\|$ can in principle be any matrix norm. It is convenient to use the Frobenius norm $\|\cdot\|_F$ because of the fact that $\|A\|_F = \|\text{vec}(A)\|_2$, where $\text{vec}(\cdot)$ stacks the columns of a matrix into one long vector. This reduces the problem to finding the 2-norm of $K_f^{(1)}(A)$, the *Kronecker form* of the Fréchet derivative: The Kronecker form is the matrix of size $n^2 \times n^2$ that satisfies

$$\text{vec}(L_f^{(1)}(A, E)) = K_f^{(1)}(A)\text{vec}(E) \text{ for all } E \in \mathbb{K}^{n \times n}.$$

Using this definition, the condition number in the Frobenius norm can be rewritten as

$$\text{cond}^{[1]}(f, A) = \max_{E \neq 0} \frac{\|L_f^{(1)}(A, E)\|_F}{\|E\|_F} = \max_{E \neq 0} \frac{\|\text{vec}(L_f^{(1)}(A, E))\|_2}{\|\text{vec}(E)\|_2} = \max_{E \neq 0} \frac{\|K_f^{(1)}(A)\text{vec}(E)\|_2}{\|\text{vec}(E)\|_2} = \|K_f^{(1)}(A)\|_2.$$

Higher-order derivatives of matrix functions are defined similarly, and in particular the second Fréchet derivative $L_f^{(2)}(A, E_1, E_2)$ is a multilinear function which satisfies

$$(2.2) \quad \|L_f^{(1)}(A + E_2, E_1) - L_f^{(1)}(A, E_1) - L_f^{(2)}(A, E_1, E_2)\| = o(\|E_2\|).$$

Sufficient conditions for the existence of the second Fréchet derivative are given by Higham and Relton [22, Theorem 3.5]. This theorem also states that $L_f^{(2)}(A, E_1, E_2)$ is equal to the upper right $n \times n$ block of $f(X_2)$,

$$(2.3) \quad L_f^{(2)}(A, E_1, E_2) = [f(X_2)]_{1:n, 1:n},$$

where X_2 is built using the following recursion

$$\begin{aligned} X_0 &= A, \\ X_k &= I_2 \otimes X_{k-1} + \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \otimes I_{2^{k-1}} \otimes E_k, \quad k \geq 1, \end{aligned}$$

with I_n the $n \times n$ identity matrix and \otimes the Kronecker product.

Using these building blocks, one can define higher-order condition numbers of matrix functions, which capture the sensitivity of the condition number itself to perturbations in A . We outline the basic approach here, following largely the derivation in [22]. We begin by defining the absolute level-2 condition number via the following equation,

$$(2.4) \quad \text{cond}^{[2]}(f, A) = \lim_{\epsilon \rightarrow 0} \sup_{\|Z\| \leq \epsilon} \frac{|\text{cond}^{[1]}(f, A + Z) - \text{cond}^{[1]}(f, A)|}{\epsilon}.$$

We can derive an upper bound for the level-2 condition number (see [22, Section 5]) by combining equations (1.1) and (2.2), which yields

$$\text{cond}^{[1]}(f, A + Z) = \max_{\|E\|=1} \|L_f^{(1)}(A, E) + L_f^{(2)}(A, E, Z) + o(\|Z\|)\|.$$

From the triangular inequality, we get the upper bound

$$\begin{aligned} |\text{cond}^{[1]}(f, A + Z) - \text{cond}^{[1]}(f, A)| &= \left| \max_{\|E\|=1} \|L_f^{(1)}(A, E) + L_f^{(2)}(A, E, Z) + o(\|Z\|)\| - \max_{\|E\|=1} \|L_f^{(1)}(A, E)\| \right| \\ &\leq \max_{\|E\|=1} \|L_f^{(2)}(A, E, Z) + o(\|Z\|)\|. \end{aligned}$$

Using this within the definition of the level-2 condition number (2.4) yields

$$\begin{aligned} \text{cond}^{[2]}(f, A) &\leq \lim_{\epsilon \rightarrow 0} \sup_{\|Z\| \leq \epsilon} \max_{\|E\|=1} \|L_f^{(2)}(A, E, Z/\epsilon) + o(\|Z\|)/\epsilon\| \\ &= \sup_{\|Z\| \leq 1} \max_{\|E\|=1} \|L_f^{(2)}(A, E, Z)\| \\ (2.5) \quad &= \max_{\|Z\|=1} \max_{\|E\|=1} \|L_f^{(2)}(A, E, Z)\|, \end{aligned}$$

where the supremum can be replaced with a maximum in a finite dimensional vector space and since $L_f^{(2)}(A, E, Z)$ is linear in Z we can have $\|Z\| = 1$. Furthermore, as $L_f^{(2)}(A, E, Z)$ is also linear in E , there exists a matrix $K_f^{(1)}(A, Z) \in \mathbb{C}^{n^2 \times n^2}$ such that

$$\text{vec}(L_f^{(2)}(A, E, Z)) = K_f^{(1)}(A, Z)\text{vec}(E).$$

Using these facts and taking the Frobenius norm in the upper bound (2.5) leads to

$$\begin{aligned} \text{cond}^{[2]}(f, A) &\leq \max_{\|Z\|_F=1} \max_{\|E\|_F=1} \|L_f^{(2)}(A, E, Z)\|_F \\ &= \max_{\|Z\|_F=1} \max_{\|\text{vec}(E)\|_2=1} \|K_f^{(1)}(A, Z)\text{vec}(E)\|_2 \\ &= \max_{\|Z\|_F=1} \|K_f^{(1)}(A, Z)\|_2 \\ &\leq \max_{\|Z\|_F=1} \|K_f^{(1)}(A, Z)\|_F \\ &= \max_{\|\text{vec}(Z)\|_2=1} \|K_f^{(2)}(A)\text{vec}(Z)\|_2 \\ (2.6) \qquad &= \|K_f^{(2)}(A)\|_2, \end{aligned}$$

where $K_f^{(2)}(A)$ is the Kronecker form of the second Fréchet derivative.

We stress that—in contrast to the level-1 condition number—the spectral norm of the Kronecker form only gives an upper bound for the level-2 condition number instead of its precise value. The exact value of the level-2 condition number can only be determined in certain special cases. Higham and Relton find formulas for the exact level-2 condition number for the matrix inverse [22, Section 5.2] and the exponential of normal matrices [22, Section 5.1], while Schweitzer finds an exact formula for certain functions of Hermitian matrices which have a strictly monotonic derivative [30, Section 3].

2.2. The structured case. In this subsection, we give a brief recap of the framework within which structured perturbations can be analyzed and proceed to derive similar results to the above for structured matrices.

Let \mathbb{S} be a smooth square matrix manifold $\mathbb{S} \subseteq \mathbb{K}^{n \times n}$ for a field \mathbb{K} . For a smooth manifold \mathbb{S} , the element $E \in \mathbb{K}^{n \times n}$ is a tangent vector of \mathbb{S} at $A \in \mathbb{S}$ if there is a smooth curve $\gamma : \mathbb{K} \rightarrow \mathbb{S}$ such that $\gamma(0) = A$, $\gamma'(0) = E$. The set

$$(2.7) \qquad T_A\mathbb{S} := \{E \in \mathbb{K}^{n \times n} \mid \exists \gamma : \mathbb{K} \rightarrow \mathbb{S} \text{ smooth with } \gamma(0) = A, \gamma'(0) = E\},$$

is called the tangent space of \mathbb{S} at A . Furthermore, we define a scalar product, from $\mathbb{K}^n \times \mathbb{K}^n$ to \mathbb{K} : $(x, y) \mapsto \langle x, y \rangle_M$ for any nonsingular matrix $M \in \mathbb{K}^{n \times n}$ by,

$$\langle x, y \rangle_M = \begin{cases} x^T M y, & \text{for real or complex bilinear forms,} \\ x^* M y, & \text{for sesquilinear forms.} \end{cases}$$

Finally, for any matrix $A \in \mathbb{K}^{n \times n}$, there exists a unique $A^* \in \mathbb{K}^{n \times n}$, called the adjoint of A with respect to the scalar product $\langle \cdot, \cdot \rangle_M$, which is given by

Table 1: Choices for M leading to well-known structures when constructing the Jordan algebra, Lie algebra, or automorphism groups \mathbb{J} , \mathbb{L} , and \mathbb{G} .

Real/Complex Bilinear forms			
M	Automorphism Group (\mathbb{G})	Jordan Algebra (\mathbb{J})	Lie Algebra (\mathbb{L})
I_n	Real orthogonals	Symmetrics	Skew-symmetrics
I_n	Complex orthogonals	Complex symmetrics	Complex skew-symmetrics
$\Sigma_{p,q}$	Pseudo-orthogonals	Pseudo-symmetrics	Pseudo skew-symmetrics
$\Sigma_{p,q}$	Complex pseudo-orthogonals	Complex pseudo-symmetrics	Complex pseudo-skew-symmetrics
R	Real perplectics	Persymmetric	Perskew-symmetrics
J	Real symplectics	Skew-Hamiltonians	Hamiltonians
J	Complex symplectics	J -skew-symmetric	J -symmetrics

$$A^\star = \begin{cases} M^{-1}A^T M, & \text{for real or complex bilinear forms,} \\ M^{-1}A^* M, & \text{for sesquilinear forms.} \end{cases}$$

We are interested in three classes of structured matrices associated with $\langle \cdot, \cdot \rangle_M$: a Jordan algebra \mathbb{J} , a Lie algebra \mathbb{L} , and an automorphism group \mathbb{G} defined by

$$\mathbb{J} := \{A \in \mathbb{K}^{n \times n} \mid A^\star = A\}, \quad \mathbb{L} := \{A \in \mathbb{K}^{n \times n} \mid A^\star = -A\}, \quad \mathbb{G} := \{A \in \mathbb{K}^{n \times n} \mid A^\star = A^{-1}\},$$

respectively. Practically important choices of M are given by the three matrices

$$(2.8) \quad \Sigma_{p,q} = \begin{bmatrix} I_p & 0 \\ 0 & -I_q \end{bmatrix}, \quad (\text{where } p + q = n), \quad R = \begin{bmatrix} & & 1 \\ & \ddots & \\ 1 & & \end{bmatrix}, \quad \text{and} \quad J = \begin{bmatrix} 0 & I_{n/2} \\ -I_{n/2} & 0 \end{bmatrix},$$

as \mathbb{J} , \mathbb{L} , and \mathbb{G} then correspond to some commonly encountered matrix structures; cf. Table 1. Note that for all choices of M from (2.8), we trivially have that $M = \mu M^\star$ with $\mu \in \{+1, -1\}$.

With these fundamentals, we are ready to define the structured level-1 condition number using differentials between tangent spaces to smooth square matrix manifolds, closely following [7].

DEFINITION 2.1. Let $f : \mathbb{S}_{\mathcal{M}} \rightarrow \mathbb{S}_{\mathcal{N}}$ be a smooth matrix function between two smooth square matrix manifolds $\mathbb{S}_{\mathcal{M}} \subseteq \mathbb{K}^{n \times n}$ and $\mathbb{S}_{\mathcal{N}} \subseteq \mathbb{K}^{n \times n}$. Let $A \in \mathbb{S}_{\mathcal{M}}$ be such that $f(A) \in \mathbb{S}_{\mathcal{N}}$ is defined. Then the absolute structured level-1 condition number of $f(A)$ is

$$(2.9) \quad \text{cond}_{\text{struc}}^{[1]}(f, A) = \lim_{\epsilon \rightarrow 0} \sup_{\substack{\|Y-A\| \leq \epsilon \\ Y \in \mathbb{S}_{\mathcal{M}}}} \frac{\|f(Y) - f(A)\|}{\epsilon}.$$

The enforced condition $Y \in \mathbb{S}_{\mathcal{M}}$ restricts the choice of the perturbation $Y - A$ to a smaller set than in the unstructured case, so from the definition of the supremum, we get

$$\text{cond}_{\text{struc}}^{[1]}(f, A) \leq \text{cond}^{[1]}(f, A).$$

The structured condition number can equivalently be expressed in terms of the *differential* of f . Precisely, if $f : \mathbb{S}_{\mathcal{M}} \rightarrow \mathbb{S}_{\mathcal{N}}$ is \mathbb{K} -differentiable, then the differential of f at the point A is the map

$$df_A : T_A \mathbb{S}_{\mathcal{M}} \rightarrow T_{f(A)} \mathbb{S}_{\mathcal{N}}, \quad df_A(\gamma'(0)) = (f \circ \gamma)'(0),$$

where $T_A \mathbb{S}_{\mathcal{M}}$ and $T_{f(A)} \mathbb{S}_{\mathcal{N}}$ are tangent spaces as defined in (2.7). Using this definition, it holds

$$(2.10) \quad \text{cond}_{\text{struc}}^{[1]}(f, A) = \|df_A\|,$$

where

$$\|df_A\| := \max_{\substack{E \in T_A \mathbb{S}_{\mathcal{M}} \\ E \neq 0}} \frac{\|df_A(E)\|}{\|E\|}.$$

The equality (2.10) is proven in [7, Theorem 2.3].

For practically computing the structured level-1 condition number, one constructs a structured equivalent of the Kronecker form: Let the columns of $B_{\mathcal{M}}$ span the tangent space $T_A \mathbb{S}_{\mathcal{M}}$ and let $y \in \mathbb{K}^p$ with $p = \dim_{\mathbb{K}} T_A \mathbb{S}_{\mathcal{M}}$. Then for $E \in T_A \mathbb{S}_{\mathcal{M}}$ we have $\text{vec}(E) = B_{\mathcal{M}}y$. Focusing on the Frobenius norm as before, the structured level-1 condition number can be expressed as

$$\text{cond}_{\text{struc}}^{[1]}(f, A) = \max_{\substack{E \in T_A \mathbb{S}_{\mathcal{M}} \\ E \neq 0}} \frac{\|\text{vec}(L_f(A, E))\|_2}{\|\text{vec}(E)\|_2} = \max_{\substack{y \in \mathbb{K}^p \\ y \neq 0}} \frac{\|K_f(A)B_{\mathcal{M}}y\|_2}{\|B_{\mathcal{M}}y\|_2} = \|K_f(A)B_{\mathcal{M}}B_{\mathcal{M}}^+\|_2,$$

where $B_{\mathcal{M}}^+$ denotes the Moore–Penrose pseudoinverse of $B_{\mathcal{M}}$. The construction of the tangent space basis $B_{\mathcal{M}}$ is given in [7, Section 3.1], which shows that for $\mathbb{S}_{\mathcal{M}} \in \{\mathbb{J}, \mathbb{L}\}$ we can construct the basis of the tangent space of \mathbb{J} and \mathbb{L} as

$$(2.11) \quad B_{\mathcal{M}} = (I_n \otimes M^{-1})D_{\mathbb{S}_{\mathcal{M}}},$$

with

$$D_{\mathbb{S}_{\mathcal{M}}} = \begin{cases} [\tilde{D}_{\mu s} & \check{I}] & \text{if } \mu s = 1, \\ \tilde{D}_{\mu s} & \text{if } \mu s = -1, \end{cases}$$

where $s = 1$ if $\mathbb{S} = \mathbb{J}$, $s = -1$ if $\mathbb{S} = \mathbb{L}$ and $\tilde{D}_{\mu s} \in \mathbb{R}^{n^2 \times n(n-1)/2}$ is a matrix with the columns

$$(e_{(i-1)n+j} + \mu s e_{(j-1)n+i})/\sqrt{2}, \quad 1 \leq i < j \leq n,$$

and $\check{I} \in \mathbb{R}^{n^2 \times n}$ has the column vectors $e_{(i-1)n+i}$, $i = 1, \dots, n$.

For an automorphism group the basis of the tangent space to $\mathbb{S}_{\mathcal{M}} = \mathbb{G}$ at A is given by

$$(2.12) \quad B_{\mathcal{M}} = (I_n \otimes AM^{-1})D_{\mathbb{L}_{\mathcal{M}}};$$

cf. [7, Section 3.2].

3. New results. In this section, we derive our main result, Lemma 3.1, and use it to find upper bounds for structured level-2 condition numbers for various different matrix structures in Sections 3.1 and 3.2. In Section 3.3, we discuss two particular cases in which it is possible to exactly compute the structured level-2 condition number.

3.1. The level-2 condition number of matrices from Lie and Jordan algebras or automorphism groups. For a smooth matrix function $f : \mathbb{S}_{\mathcal{M}} \rightarrow \mathbb{S}_{\mathcal{N}}$, we can define the structured level-2 condition number similarly to (2.4), but restricting the perturbation matrix Z to $\mathbb{S}_{\mathcal{M}}$,

$$(3.13) \quad \text{cond}_{\text{struc}}^{[2]}(f, A) = \lim_{\epsilon \rightarrow 0} \sup_{\substack{A+Z \in \mathbb{S}_{\mathcal{M}} \\ \|Z\| \leq \epsilon}} \frac{|\text{cond}_{\text{struc}}^{[1]}(f, A+Z) - \text{cond}_{\text{struc}}^{[1]}(f, A)|}{\epsilon}.$$

The following lemma provides an upper bound of the structured level-2 condition number of $f(A)$.

LEMMA 3.1. *Let $f : \mathbb{S}_{\mathcal{M}} \rightarrow \mathbb{S}_{\mathcal{N}}$ be a smooth matrix function between two smooth square matrix manifolds $\mathbb{S}_{\mathcal{M}} \subseteq \mathbb{K}^{n \times n}$ and $\mathbb{S}_{\mathcal{N}} \subseteq \mathbb{K}^{n \times n}$. Let $A \in \mathbb{S}_{\mathcal{M}}$ be such that $f(A) \in \mathbb{S}_{\mathcal{N}}$ is defined. Then an upper bound for the structured level-2 condition number of $f(A)$ is given by*

$$(3.14) \quad \text{cond}_{\text{struc}}^{[2]}(f, A) \leq \|(B_{\mathcal{M}}B_{\mathcal{M}}^+ \otimes I_{n^2})K_f^{(2)}(A)B_{\mathcal{M}}B_{\mathcal{M}}^+\|_2,$$

where the columns of $B_{\mathcal{M}}$ span $T_A\mathbb{S}_{\mathcal{M}}$, the tangent space of $\mathbb{S}_{\mathcal{M}}$ at A .

Proof. Using the first and the second Fréchet derivatives gives

$$\text{cond}_{\text{struc}}^{[1]}(f, A+Z) = \max_{\substack{E \in \mathbb{S}_{\mathcal{M}}, Z \in \mathbb{S}_{\mathcal{M}}, \\ \|E\|=1}} \|L_f^{(1)}(A, E) + L_f^{(2)}(A, E, Z) + o(\|Z\|)\|.$$

Enforcing the structure on the perturbation matrices E and Z with $E, Z \in \mathbb{S}_{\mathcal{M}}$ yields the following upper bound of the structured level-2 condition number:

$$(3.15) \quad \begin{aligned} \text{cond}_{\text{struc}}^{[2]}(f, A) &\leq \max_{\substack{Z \in T_A\mathbb{S}_{\mathcal{M}} \\ \|Z\|_F=1}} \max_{\substack{E \in T_A\mathbb{S}_{\mathcal{M}} \\ \|E\|_F=1}} \|L_f^{(2)}(A, E, Z)\|_F \\ &= \max_{\substack{Z \in T_A\mathbb{S}_{\mathcal{M}} \\ \|Z\|_F=1}} \max_{\substack{E \in T_A\mathbb{S}_{\mathcal{M}} \\ \|\text{vec}(E)\|_2=1}} \|K_f^{(1)}(A, Z)\text{vec}(E)\|_2 \\ &= \max_{\substack{Z \in T_A\mathbb{S}_{\mathcal{M}} \\ \|Z\|_F=1}} \max_{\substack{y_1 \in \mathbb{K}^p \\ y_1 \neq 0}} \|K_f^{(1)}(A, Z)B_{\mathcal{M}}y_1\|_2 \\ &= \max_{\substack{Z \in T_A\mathbb{S}_{\mathcal{M}} \\ \|Z\|_F=1}} \|K_f^{(1)}(A, Z)B_{\mathcal{M}}B_{\mathcal{M}}^+\|_2, \end{aligned}$$

where, for the last equality, we used the fact that $B_{\mathcal{M}} = B_{\mathcal{M}}B_{\mathcal{M}}^+B_{\mathcal{M}}$ together with the definition of the 2-norm,

$$\|K_f^{(1)}(A, Z)B_{\mathcal{M}}B_{\mathcal{M}}^+\|_2 = \max_{\substack{y_1 \in \mathbb{K}^p \\ y_1 \neq 0}} \frac{\|K_f^{(1)}(A, Z)B_{\mathcal{M}}B_{\mathcal{M}}^+B_{\mathcal{M}}y_1\|_2}{\|B_{\mathcal{M}}y_1\|_2}.$$

Bounding the 2-norm by the Frobenius norm, we further obtain, starting from (3.15),

$$\begin{aligned}
 \text{cond}_{\text{struc}}^{[2]}(f, A) &\leq \max_{\substack{Z \in T_A \mathbb{S}_{\mathcal{M}} \\ \|Z\|_F=1}} \|K_f^{(1)}(A, Z)B_{\mathcal{M}}B_{\mathcal{M}}^+\|_F \\
 &= \max_{\substack{Z \in T_A \mathbb{S}_{\mathcal{M}} \\ \|\text{vec}(Z)\|_2=1}} \|(B_{\mathcal{M}}B_{\mathcal{M}}^+ \otimes I_{n^2})K_f^{(2)}(A)\text{vec}(Z)\|_2 \\
 &= \max_{\substack{y_2 \in \mathbb{K}^p \\ y_2 \neq 0}} \|(B_{\mathcal{M}}B_{\mathcal{M}}^+ \otimes I_{n^2})K_f^{(2)}(A)B_{\mathcal{M}}y_2\|_2 \\
 &= \|(B_{\mathcal{M}}B_{\mathcal{M}}^+ \otimes I_{n^2})K_f^{(2)}(A)B_{\mathcal{M}}B_{\mathcal{M}}^+\|_2,
 \end{aligned}$$

again using $B_{\mathcal{M}} = B_{\mathcal{M}}B_{\mathcal{M}}^+B_{\mathcal{M}}$. □

We now modify [22, Algorithm 4.2] to obtain an upper bound for the structured level-2 condition number. It requires any algorithm to compute $L_f(A, E_1, E_2)$ for $f : \mathbb{S}_{\mathcal{M}} \rightarrow \mathbb{S}_{\mathcal{N}}$, where $\mathbb{S}_{\mathcal{M}} \subseteq \mathbb{K}^{n \times n}$ is a smooth matrix manifold and $E_i \in T_A \mathbb{S}_{\mathcal{M}}$, i.e., $\text{vec}(E_i) = B_{\mathcal{M}}y_i$ for some $y_i \in \mathbb{K}^p$, and returns $\gamma = \|(B_{\mathcal{M}}B_{\mathcal{M}}^+ \otimes I_{n^2})K_f^{(2)}(A)B_{\mathcal{M}}B_{\mathcal{M}}^+\|_2$. The resulting method is described in Algorithm 1.

Algorithm 1 Upper bound for structured level-2 condition number

Input: Matrix $A \in \mathbb{R}^{n \times n}$, function f , basis $B_{\mathcal{M}}$ of $T_A \mathbb{S}_{\mathcal{M}}$

Output: Upper bound $\gamma = \|(B_{\mathcal{M}}B_{\mathcal{M}}^+ \otimes I_{n^2})K_f^{(2)}(A)B_{\mathcal{M}}B_{\mathcal{M}}^+\|_2$ for structured level-2 condition number

- 1: $\Psi \leftarrow 0_{n^2 \times p}$
 - 2: $\Omega \leftarrow 0_{n^4 \times p}$
 - 3: **for** $i = 1 : p$ **do**
 - 4: Choose $E_1 \in \mathbb{K}^{n \times n}$ such that $\text{vec}(E_1) = B_{\mathcal{M}}e_i$
 - 5: **for** $j = 1 : p$ **do**
 - 6: Choose $E_2 \in \mathbb{K}^{n \times n}$ such that $\text{vec}(E_2) = B_{\mathcal{M}}e_j$
 - 7: Compute $L = L_f^{(2)}(A, E_1, E_2)$ using the relation (2.3)
 - 8: $\Psi e_j \leftarrow \text{vec}(L)$
 - 9: **end for**
 - 10: $\Omega e_i \leftarrow \text{vec}(\Psi)$
 - 11: **end for**
 - 12: $\gamma = \|\Omega\|_2$
-

We briefly comment on the computational cost of Algorithm 1: For typical practically relevant matrix functions, algorithms are available for evaluating them at a cost of $\mathcal{O}(n^3)$. These also directly transfer (via the relation (2.3)) to algorithms for computing $L_f(A, E_1, E_2)$ at cost $\mathcal{O}(n^3)$, albeit with an additional constant $4^3 = 64$ hidden in the \mathcal{O} , as the matrix X_2 in (2.3) is of size $4n \times 4n$. Thus, the overall cost of the algorithm is $\mathcal{O}(p^2n^3)$, where p is the dimension of the tangent space $T_A \mathbb{S}_{\mathcal{M}}$ and the constant can be expected to be quite large. This makes the algorithm prohibitively expensive even for medium-scale problems (a typical limitation that is also already present when computing unstructured level-2 or even level-1 condition numbers).

The computational cost of individual Fréchet derivative evaluations can be reduced somewhat by using complex-step approximations [2] or quadrature-based algorithms [30]. While in general, when A and the E_i

are unstructured, this does not change the asymptotic scaling of Algorithm 1, experiments in [30] indicate that run time can be reduced by about one order of magnitude. If the direction terms E_i are rank-one matrices (which, e.g., happens if the columns of $B_{\mathcal{M}}$ are certain canonical unit vectors), then run time of Algorithm 1 can potentially be reduced to $\mathcal{O}(p^2n^2)$ if linear systems with A can be solved efficiently; see [30, Section 4.3] for details.

3.2. The level-2 condition number of (quasi-)triangular matrices. Another type of matrix structure that differs from those considered in Sections 2.2 and 3.1 is (quasi-)triangularity, which plays an important role in many algorithms for computing matrix functions. An upper quasi-triangular matrix is a block upper triangular matrix with diagonal blocks of size at most 2×2 . These commonly occur when using the Schur decomposition: For general A , a *Schur decomposition* $A = QUQ^*$ with Q unitary and U upper triangular is always guaranteed to exist. Note, however, that U and Q may be complex even if A is real. In the case of real A , one can instead use a *real Schur decomposition* $A = QUQ^T$, where the matrices U and Q are both real, but U is now upper quasi-triangular.

Given a (real) Schur decomposition, the relation $f(A) = Qf(U)Q^*$ shows that evaluating $f(A)$ can essentially be reduced to the evaluating f at a (quasi-)triangular matrix. This forms the basis of many numerical methods for evaluating $f(A)$ at small and dense matrices A : the Schur–Parlett algorithm [11, 21], for example.

Structured level-1 condition numbers of $f(U)$ have recently been investigated in [1], and we extend this concept to the level-2 case. Quasi-triangular matrices do not fit into the framework of Lie and Jordan algebras or automorphism groups induced by a scalar product and have an arguably simpler structure. We recall some important properties:

1. The set \mathbb{S} of quasi-triangular matrices (with a fixed structure of diagonal blocks) forms a subalgebra of $\mathbb{K}^{n \times n}$ and is thus in particular a linear space and a smooth matrix manifold.
2. As \mathbb{S} is a linear space, we have $T_U\mathbb{S} = \mathbb{S}$ for any $U \in \mathbb{S}$.
3. As any matrix function $f(U)$ is a polynomial in U and \mathbb{S} is a subalgebra of $\mathbb{K}^{n \times n}$, we have $f : \mathbb{S} \rightarrow \mathbb{S}$, i.e., a function of a quasi-triangular matrix is a quasi-triangular matrix with the same diagonal block structure.

In light of the above observations, Lemma 3.1 also applies to the case of quasi-triangular matrices, and for $B_{\mathcal{M}}$ we can take a basis of the space of upper quasi-triangular matrices (with the induced block structure). Such a basis is easily constructed from canonical unit vectors: Using the notation introduced in [1], let $d \in \mathbb{R}^{n-1}$ denote a vector that encodes the diagonal block structure of U , i.e.,

$$d_i = \begin{cases} 1, & \text{if } u_{i+1,i} \neq 0 \\ 0, & \text{otherwise,} \end{cases} \quad \text{for } i = 1, \dots, n-1.$$

To clearly indicate that the block structure is fixed, we denote the algebra of all quasi-triangular matrices with block structure encoded by d as \mathbb{S}_d . Clearly, $\dim(\mathbb{S}_d) = \frac{n(n+1)}{2} + \text{nnz}(d)$, where $\text{nnz}(d)$ denotes the number of nonzero entries. A (vectorized) basis of \mathbb{S}_d is given by

$$(3.16) \quad \{e_j \otimes e_i : j = 1, \dots, n, i = 1, \dots, j\} \cup \{e_j \otimes e_{j+1} : j = 1, \dots, n-1 \text{ with } d_j = 1\}.$$

The basis (3.16) is orthonormal by construction, so if we take these basis vectors as columns of $B_{\mathcal{M}}$, we have $B_{\mathcal{M}}^+ = B_{\mathcal{M}}^*$.

3.3. Exact structured level-2 condition numbers for special cases. So far, we were only able to determine upper bounds for the structured level-2 condition number, which is in general all one can hope to obtain. However, for a few particular combinations of matrix class, function f and used norm, it is possible to give an exact, closed form for $\text{cond}_{\text{struc}}^{[2]}(f, A)$.

Exponential of skew-Hermitian matrices. When using the spectral norm (instead of the Frobenius norm that is typically employed), one can exactly compute the structured level-2 condition number when A comes from the Lie Algebra of skew-Hermitian matrices. Note that exponentials of skew-Hermitian matrices, e.g., play an important role in the simulation of quantum systems, where according to the Schrödinger equation, the evolution of states is governed by $\exp(iH)$, where H is the (Hermitian) Hamiltonian of the system; see, e.g., [27]. As a special case of the very general result that the exponential map takes a Lie algebra into its corresponding Lie group, it is well known that the exponential of a skew-Hermitian matrix is unitary, a property that is fundamental for proving the following result.

PROPOSITION 3.2. *Let $A \in \mathbb{C}^{n \times n}$ be skew-Hermitian. Then, in the spectral norm,*

$$\text{cond}_{\text{struc}}^{[2]}(\exp, A) = 0.$$

Proof. By a result from [32, Section 4] for the unstructured condition number, it holds for any normal matrix M that in the spectral norm $\text{cond}^{[1]}(\exp, M) = e^{\alpha(M)}$, where $\alpha(M)$ denotes the largest real part of any eigenvalue of M (the so-called *spectral abscissa*). Any skew-Hermitian matrix A is normal and has eigenvalues on the imaginary axis, so that $\alpha(A) = 0$, and therefore $\text{cond}^{[1]}(\exp, A) = e^0 = 1$. We now argue that the same holds for the structured level-1 condition number, i.e., $\text{cond}_{\text{struc}}^{[1]}(\exp, A) = 1$. Consider the skew-Hermitian matrix $Y_{\varepsilon} = A + i\varepsilon I$. We then have $\exp(Y_{\varepsilon}) = e^{i\varepsilon} \exp(A)$ and $\|Y_{\varepsilon} - A\| = \varepsilon$. Taking this into account, we find

$$(3.17) \quad \lim_{\varepsilon \rightarrow 0} \frac{\|\exp(Y_{\varepsilon}) - \exp(A)\|}{\varepsilon} = \lim_{\varepsilon \rightarrow 0} \frac{\|(e^{i\varepsilon} - 1)\exp(A)\|}{\varepsilon} = \lim_{\varepsilon \rightarrow 0} \frac{|e^{i\varepsilon} - 1| \cdot \|\exp(A)\|}{\varepsilon} = \lim_{\varepsilon \rightarrow 0} \frac{|e^{i\varepsilon} - 1|}{\varepsilon} = 1,$$

where we have used in the second to last equality that $\exp(A)$ is unitary and therefore has spectral norm equal to one. From (3.17), it follows that the supremum in (2.9) is at least one. On the other hand, the structured condition number is always bounded above by the unstructured one, so they must agree. Using this, the assertion of the proposition directly follows: From the definition (3.13), we have

$$\text{cond}_{\text{struc}}^{[2]}(\exp, A) = \lim_{\varepsilon \rightarrow 0} \sup_{\substack{A+Z \in \mathbb{S}_{\mathcal{M}} \\ \|Z\| \leq \varepsilon}} \frac{|\text{cond}_{\text{struc}}^{[1]}(\exp, A+Z) - \text{cond}_{\text{struc}}^{[1]}(\exp, A)|}{\varepsilon} = \lim_{\varepsilon \rightarrow 0} \sup_{\substack{A+Z \in \mathbb{S}_{\mathcal{M}} \\ \|Z\| \leq \varepsilon}} \frac{|1 - 1|}{\varepsilon} = 0,$$

where we have used that $A + Z$ is again skew-Hermitian. □

The given proposition states that for a skew-Hermitian matrix A when we analyze its matrix exponential $\exp(A)$ using the spectral norm, the structured level-2 condition number is remarkably 0. This result implies that the structured level-1 condition number is constant and remains unchanged by any perturbation, as long as the perturbing matrix maintains the skew-Hermitian structure.

Exponential of Hermitian matrices. Next, we investigate the Hermitian case. As in the skew-Hermitian case, the structured and unstructured level-1 condition numbers agree, and this time, this carries over to the level-2 condition numbers.

PROPOSITION 3.3. *Let $A \in \mathbb{C}^{n \times n}$ be Hermitian. Then, in the spectral norm,*

$$\text{cond}_{\text{struc}}^{[2]}(\exp, A) = \text{cond}^{[2]}(\exp, A) = e^{\lambda_{\max}},$$

where λ_{\max} denotes the largest eigenvalue of A .

Proof. Proceeding as in the proof of Proposition 3.2, we have that $\text{cond}_{\text{struc}}^{[1]}(\exp, A) = \text{cond}^{[1]}(\exp, A) = e^{\alpha(A)} = e^{\lambda_{\max}}$. According to [22, Theorem 5.2], it also holds that $\text{cond}^{[2]}(\exp, A) = e^{\lambda_{\max}}$. Clearly, for any Hermitian Y with $\|Y - A\| \leq \varepsilon$ we have $e^{\alpha(Y)} \leq e^{\lambda_{\max} + \varepsilon}$, as the spectral norm of a Hermitian matrix is simply its largest eigenvalue. The upper bound is attained for the Hermitian matrix $Y = A + \varepsilon I$, and it follows from the definition of the structured level-2 condition number that

$$\text{cond}_{\text{struc}}^{[2]}(\exp, A) = \lim_{\varepsilon \rightarrow 0} \frac{|e^{\lambda_{\max} + \varepsilon} - e^{\lambda_{\max}}|}{\varepsilon} = e^{\lambda_{\max}}. \quad \square$$

REMARK 3.4. *One might argue that the results of Proposition 3.2 and 3.3 are of limited practical relevance. However, they highlight quite interesting properties and difficulties related to level-2 (structured) condition numbers: In both cases, the structured and unstructured level-1 condition numbers agree, but the implications for the level-2 condition number are fundamentally different. While in the Hermitian case, the structured and unstructured level-2 condition number also always agree, they can never be equal in the skew-Hermitian case (the level-2 unstructured condition number of the exponential of a skew-Hermitian matrix is always equal to 1 according to [22, Theorem 5.2]). As the structured condition number is zero, this even shows that without additional assumptions, it is impossible to derive any meaningful bounds on the ratio between the unstructured and structured condition number,*

$$\frac{\text{cond}^{[2]}(f, A)}{\text{cond}_{\text{struc}}^{[2]}(f, A)} \leq C,$$

with $C > 1$ for relating the structured and unstructured condition numbers (i.e., in general, the unstructured condition number can be “infinitely worse” than the structured one).

4. Numerical experiments. The aim of this section is to compare the structured and unstructured level-2 condition numbers for the matrix logarithm, matrix square root, and matrix exponential.

One fundamental difficulty in gauging the quality of the results obtained in Sections 3.1 and 3.2 is that there is no known algorithm to compute the exact level-2 condition number (in either the structured or unstructured case) except for very particular special cases where analytic forms are known; see the discussion at the end of Section 2.2 for the unstructured case and Section 3.3 above for the structured one. Thus, instead of comparing the structured level-2 condition number to the unstructured level-2 condition number, we can only compare upper bounds for both quantities.

This has a fundamental limitation in that we do not know the sharpness of the bounds: if an arbitrary upper bound for quantity Q_1 lies below an arbitrary upper bound for quantity Q_2 , this does not tell us anything about the relation between Q_1 and Q_2 . Thus, even when (3.14) lies orders of magnitude below (2.6), this does not necessarily mean that $\text{cond}^{[1]}(f, A)$ and $\text{cond}_{\text{struc}}^{[1]}(f, A)$ are different. Of course, one can

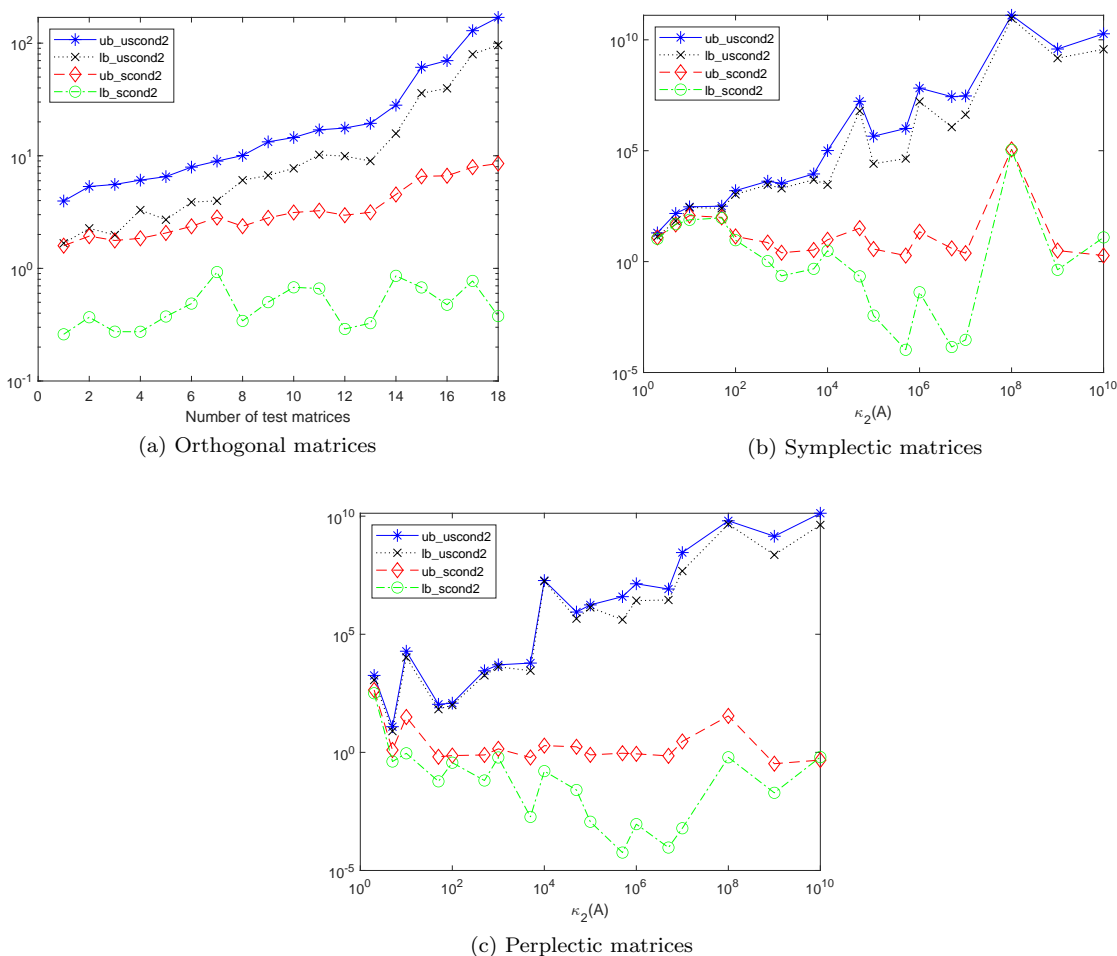


Figure 1: Structured and unstructured level-2 condition numbers for the matrix logarithm, comparing upper and lower bounds.

intuitively expect that both bounds (2.6) and (3.14) are similar in quality as they originate in strongly related concepts, but this is not guaranteed.

In order to obtain more robust results, we also compute lower bounds for the level-2 condition numbers in our experiments: If an upper bound for the structured condition number lies well below a lower bound for the unstructured condition number, then this is clearly also the case for the exact quantities, which must differ by *at least* the same amount. To obtain lower bounds, observe that in light of (2.4), for any particular matrix Z with $\|Z\|_F = 1$, we clearly have

$$(4.18) \quad \text{cond}^{[2]}(f, A) \geq \lim_{\epsilon \rightarrow 0} \frac{|\text{cond}^{[1]}(f, A + \epsilon Z) - \text{cond}^{[1]}(f, A)|}{\epsilon}.$$

Thus, the right-hand side of (4.18) gives a lower bound of the unstructured condition number for any fixed Z . Clearly, a similar relation also holds for the structured condition number, provided that Z is chosen from

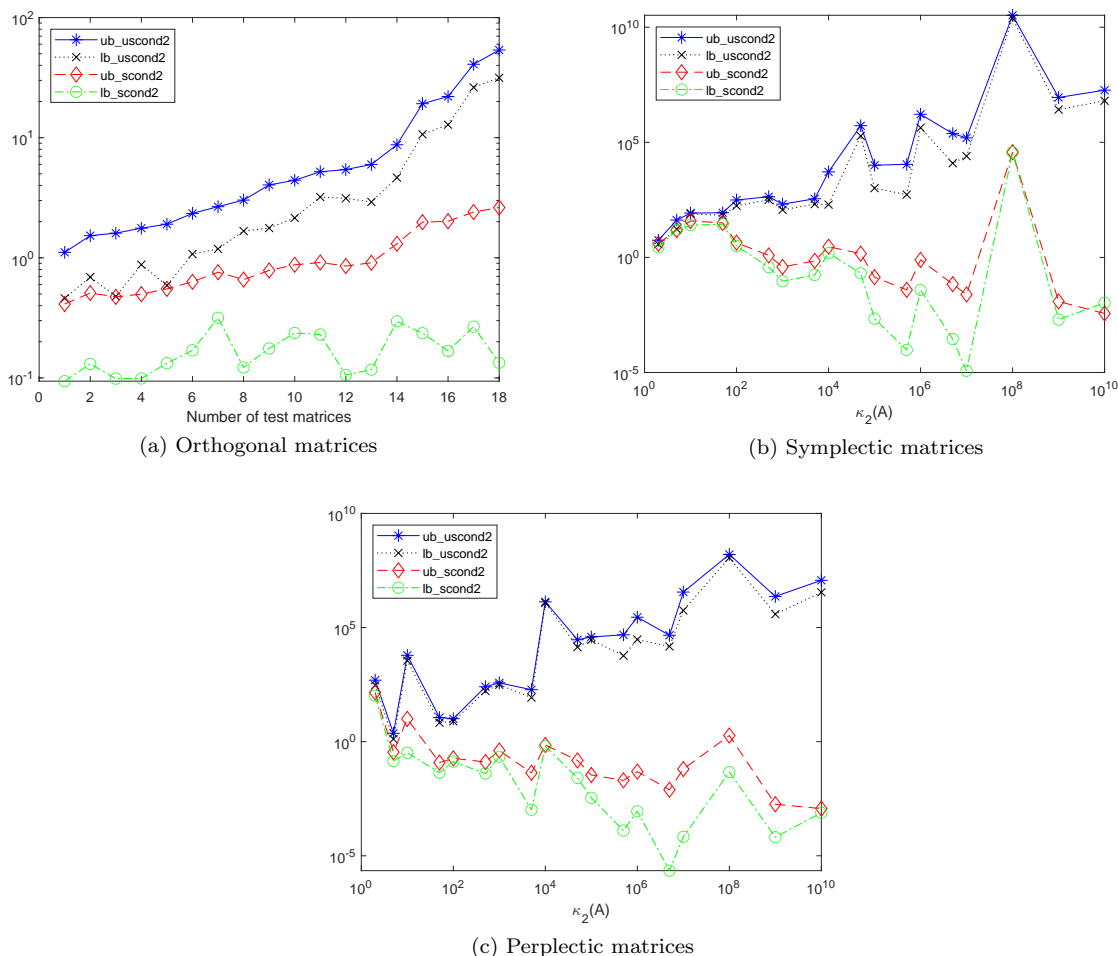


Figure 2: Structured and unstructured level-2 condition numbers for the matrix square root, comparing upper and lower bounds.

the according tangent space. To obtain lower bounds in our experiments, we use the MATLAB built-in function `fminsearch`, which implements the derivative-free simplex search method from [26] for maximizing the right-hand side of (4.18) by choosing Z . Within our numerical experiments, we replace the limit and use a finite difference quotient instead, with $\epsilon = 10^{-3}$. Strictly speaking, what we obtain this way is only an *approximate lower bound*, but a comparison to the upper bounds indicates that the result of the optimization procedure is reliable.

We compare the upper bounds from Sections 2.1, 3.1, and 3.2, as well as the approximate lower bounds (4.18). We use MATLAB R2022a on a machine with an AMD Ryzen 7 3700X 8-core CPU to run the experiments. The corresponding MATLAB code is available at https://github.com/MarcelSchweitzer/structured_level2_condition_numbers.

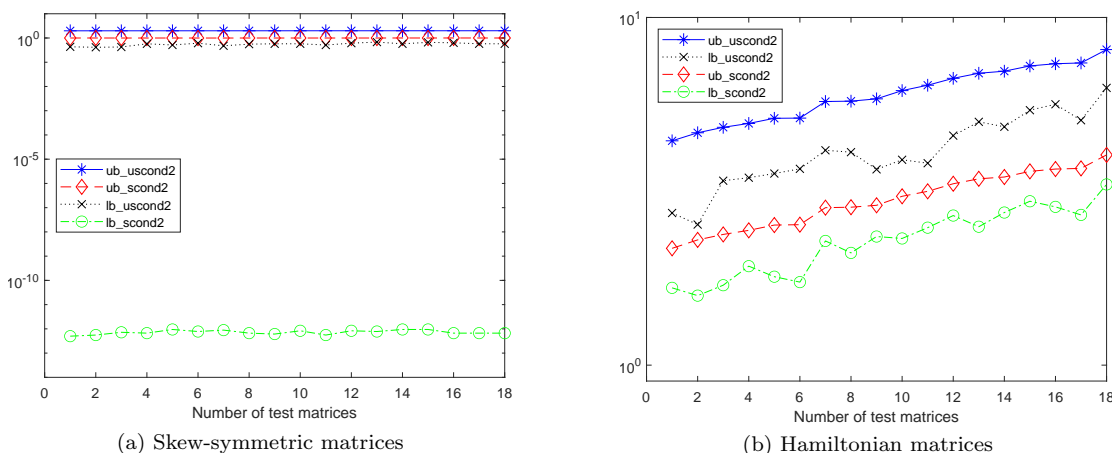


Figure 3: Structured and unstructured level-2 condition numbers for the matrix exponential, comparing upper and lower bounds.

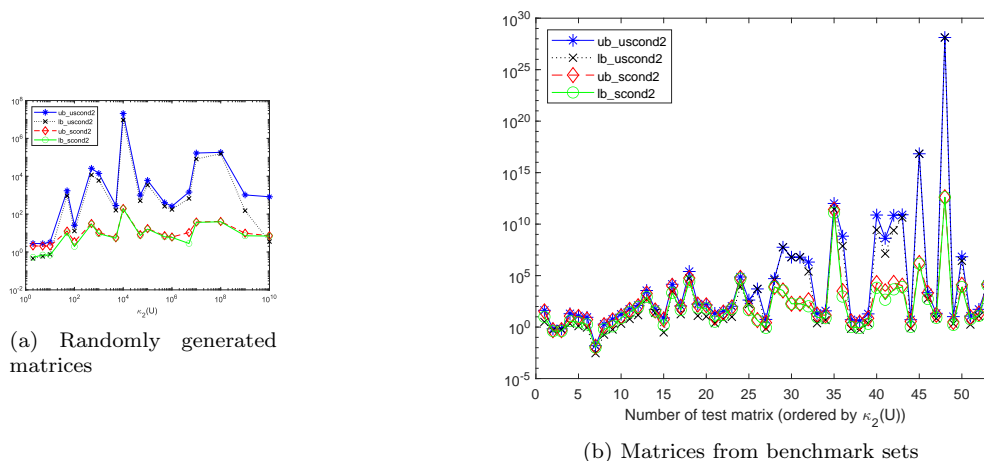


Figure 4: Structured and unstructured level-2 condition numbers for the exponential of upper quasi-triangular matrices, comparing upper and lower bounds. The condition numbers of the benchmark matrices in the right plot ranges from 2.4 to $3 \cdot 10^{31}$. The first matrix for which the structured and unstructured bounds differ significantly (number 26) has a condition number of $1.2 \cdot 10^4$

We build test matrices as follows.

- Orthogonal, symplectic, and perplectic matrices are built using Jagger’s toolbox [24].
- Hamiltonian matrices are built as

$$A = \begin{bmatrix} X & G \\ F & -X^T \end{bmatrix}, \quad \text{where } F^T = F \quad \text{and} \quad G^T = G.$$

- Quasi-triangular matrices are artificially generated as follows: We generate a diagonal matrix D with equidistantly distributed entries in the interval $[-c, -1]$ for a specified parameter c . We then

compute $A = XDX^{-1}$ with X a randomly sampled matrix and perform a real Schur decomposition $A = QUQ^T$ to obtain the upper quasi-triangular matrix U .

- Alternatively, we obtain quasi-triangular matrices by performing a (real) Schur decomposition of a large set of common benchmark matrices from the matrix function literature; see Appendix A for details on the test set.

Depending on the matrix function, we construct $B_{\mathcal{M}}$ using either equation (2.11) or (2.12).

The upper bounds for the structured and unstructured level-2 condition numbers are computed using Algorithm 1 and equation (2.6), respectively. For symplectic, perplectic, and quasi-triangular matrices, the bounds are plotted against the 2-norm condition number for matrix inversion, $\kappa_2(A) = \|A\|_2 \|A^{-1}\|_2$, whereas for orthogonal matrices we sort the test matrices by increasing values of the unstructured condition number. Within the legend of each plot, the acronyms **ub** and **lb** are used to indicate upper and lower bounds; meanwhile, **scond2** and **uscond2** denote the structured and unstructured level-2 condition numbers.

Experiment 4.1. We first consider the matrix logarithm $\log A : \mathbb{G} \rightarrow \mathbb{L}$ and by taking $\mathbb{S}_{\mathcal{M}} = \mathbb{G}$, we construct $B_{\mathcal{M}}$ using equation (2.12). In Algorithm 1, we compute the matrix logarithm by the MATLAB function `logm`, which uses the improved inverse scaling and squaring method [5]. The upper bounds for the structured and unstructured level-2 condition numbers are shown in Fig. 1 for orthogonal, symplectic, and perplectic matrices $A \in \mathbb{R}^{4 \times 4}$. Comparing the results obtained from different test matrices, we observe that in particular in the symplectic and perplectic case, the structured condition number is much smaller than the unstructured one (in fact, the structured level-2 condition number appears to stay almost constant across all considered test matrices). In the orthogonal case, the difference is less pronounced, but can still be observed clearly. Both the magnitude and the slope at which it increases are substantially smaller for the structured level-2 condition number.

The lower bounds that we also report confirm that in most cases, our upper bounds are quite tight and the structured condition number is indeed guaranteed to be much smaller than the unstructured one. Let us briefly comment on the fact that for the most ill-conditioned symplectic and perplectic matrices, the lower bound for the unstructured condition number that we report lies *above* the upper bound. This is likely caused by the fact that it is actually only an approximate lower bound (cf. the discussion at the beginning of this section) combined with the severe ill-conditioning of the matrix under consideration. Apart from this, our results (also those reported in later experiments) still indicate that in general the lower bounds can be expected to be quite reliable.

Experiment 4.2. We take the matrix square root $A^{1/2} : \mathbb{G} \rightarrow \mathbb{G}$ and $B_{\mathcal{M}}$ is built by equation (2.12). For computing the matrix square root, we use `sqrtn` which is based on the method given in [12]. The bounds are shown in Fig. 2 for orthogonal, symplectic, and perplectic matrices $A \in \mathbb{R}^{4 \times 4}$. We observe similar results as in the previous experiment. In particular, there is a dramatic increase in the unstructured level-2 condition numbers for ill-conditioned matrices, while the structured one again stays roughly constant for symplectic and perplectic matrices and only increases moderately for orthogonal matrices.

The lower bounds again confirm that our results are reliable and show that the upper bounds are quite tight, in particular in the symplectic and perplectic case.

Experiment 4.3. In this experiment, we consider the matrix exponential $e^A : \mathbb{L} \rightarrow \mathbb{G}$ for skew-symmetric and Hamiltonian matrices $A \in \mathbb{R}^{4 \times 4}$. We again construct $B_{\mathcal{M}}$ by equation (2.11) and use the MATLAB function `expm`, which is based on a scaling and squaring algorithm [3] to compute the exponential. Figure 3 demonstrates that the upper bounds for the structured and unstructured level-2 condition number behave very similarly in both cases, with the bounds for the unstructured condition number being almost exactly

two times as large as the structured one for both kinds of test matrices. In light of Proposition 3.2 (although this result uses the spectral norm and not the Frobenius norm as in this experiment), it is interesting to observe that in the skew-symmetric case, the lower bound stays in the order of 10^{-13} . This value close to machine precision appears to confirm that the exact value is zero.

Experiment 4.4. We take randomly generated (as described at the beginning of this section) upper quasi-triangular matrices $U \in \mathbb{R}^{8 \times 8}$, with the parameter c ranging between 2 and 10^{10} and $f(U) = \exp(U)$. Note that we increase the matrix size compared to the previous experiments to allow for more variety in the diagonal block structure of U . We compute the upper bounds (2.6) and (3.14) as well as lower bounds for the structured and unstructured level-2 condition number. The results are depicted on the left-hand side of Fig. 4. As it was the case for the other considered matrix structures, the bounds again confirm that structured level-2 conditioning can be much better than unstructured conditioning, in particular for the more ill-conditioned examples. The lower bounds obtained by optimization are also mostly very tight.

Experiment 4.5. In the final experiment, instead of randomly generating matrices, we compute bounds for the structured and unstructured level-2 condition numbers of $\exp(U)$, where U is taken as the upper quasi-triangular factor of the (real) Schur decomposition of a wide range of matrices from the matrix function literature; see Appendix A for details on the test set, which contains 53 matrices in total. The two-norm condition numbers of these benchmark matrices range from 2.4 to $3 \cdot 10^{31}$. All considered matrices are non-normal, so that the matrix U from their Schur decomposition is not diagonal. For matrices that are scalable in size (which is, e.g., the case for most matrices coming from built-in MATLAB functions), we choose them to be from $\mathbb{R}^{10 \times 10}$ (or the closest possible value if there are certain restrictions on the matrix size n). The resulting condition number bounds are shown on the right-hand side of Fig. 4. In particular for those test matrices which are rather well-conditioned, the upper bounds for structured and unstructured condition numbers are very similar and differ by less than one order of magnitude. However, for a large portion of the more ill-conditioned examples (starting from a two-norm condition number of $1.2 \cdot 10^4$), the upper bound of the structured condition number is often several orders of magnitude lower than that for the unstructured one. In almost all cases, the lower bounds confirm that the upper bounds are quite tight for those matrices, indicating that conclusions drawn from those bounds are indeed reasonable. Note that for a few of the more well-conditioned matrices, the lower bound for the unstructured condition number actually lies below the upper bound for the structured condition number, so that no guaranteed conclusions can be drawn. As the results suggest that both condition numbers are close to each other anyway, this is not an essential problem, though.

5. Conclusions. This work compares upper and the lower bounds of structured level-2 condition numbers of matrix functions with unstructured ones. Our results show that the difference between the bounds of the structured and unstructured level-2 condition numbers depends on the choice of the matrices and matrix functions: While for orthogonal matrices, the upper bound of the structured level-2 condition number is not significantly smaller than the upper bound of the unstructured one, the importance of structure preserving algorithms emerges for the matrix logarithm and for the matrix square root of matrices in automorphism groups. Further analysis can be done for other matrix functions and matrix factorizations. Additionally, finding more efficient algorithms for computing the level-2 condition number would be very important for making the concept practically usable in actual computations.

Acknowledgments. We are grateful to Awad Al-Mohy for providing us his test suite of quasi-triangular matrices from [1] for use in our Experiment 4.5.

Appendix A. Details on test matrices used in Experiment 4.5. In this section, we provide a comprehensive list of the matrices that have been used in Experiment 4.5, i.e., for comparing structured

and unstructured condition numbers of quasi-triangular matrices. Most of the test matrices are available via built-in MATLAB routines, but some also come from other toolboxes and benchmark collections. In general, for a test matrix A from some benchmark collection, we first compute the Schur decomposition $A = QUQ^*$ (or the real Schur decomposition $A = QUQ^T$ if A is real) and then use the upper triangular factor U as input for the condition number estimation algorithms.

Table 2: Test matrices used in Experiment 4.5: (a) Matrices from MATLAB gallery, (b) other built-in MATLAB matrices, (c) matrices from the Matrix Computation Toolbox [18]. *: The entries of these matrices were scaled by $\frac{10}{\|A\|_1}$ to prevent overflow in the matrix exponential.

(a)				(b)	(c)
binomial*	forsythe	krylov*	rando	rand	gfpp
chebspec	frank	leslie	randsvd	randn	makejcf
chebvand	gearmat	lesp	redheff	pascal*	rschur
chow	grcar	lotkin	riemann		vand
clement	invhess	parter	smoke		
cycol	invol*	randcolu			
dramadah	kahan	randjorth*			

Table 2 summarizes a large part of the matrices we used. Additionally, we used matrices that are not available via built-in MATLAB or toolbox functions:

The following two matrices are taken from [33, Test 3–4],

$$A_1 = \begin{bmatrix} -131 & 19 & 18 \\ -390 & 56 & 54 \\ -387 & 57 & 52 \end{bmatrix}, \quad A_2 = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 10^{-10} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

The next two matrices are from [15, Example 3.10, Example 4.4],

$$A_3 = \frac{1}{8} \begin{bmatrix} 1.3 & 1.3 & 1.3 & 1.3 & 10^6 & 10^6 & 10^6 & 10^6 \\ 1.3 & 1.3 & 1.3 & 1.3 & 10^6 & 10^6 & 10^6 & 10^6 \\ 1.3 & 1.3 & 1.3 & 1.3 & 10^6 & 10^6 & 10^6 & 10^6 \\ 1.3 & 1.3 & 1.3 & 1.3 & 10^6 & 10^6 & 10^6 & 10^6 \\ 0 & 0 & 0 & 0 & -1.3 & -1.3 & -1.3 & -1.3 \\ 0 & 0 & 0 & 0 & -1.3 & -1.3 & -1.3 & -1.3 \\ 0 & 0 & 0 & 0 & -1.3 & -1.3 & -1.3 & -1.3 \\ 0 & 0 & 0 & 0 & -1.3 & -1.3 & -1.3 & -1.3 \end{bmatrix}, \quad A_4 = \begin{bmatrix} \exp(0.1) & 10^6 \cdot \exp(0.1) \\ 0 & 10^{-8} + \exp(0.1) \end{bmatrix}.$$

The following matrix is from [25, Section 4],

$$A_5 = \begin{bmatrix} 48 & -49 & 50 & 49 \\ 0 & -2 & 100 & 0 \\ 0 & -1 & -2 & 1 \\ -50 & 50 & 50 & -52 \end{bmatrix}.$$

The next two matrices are adapted from [29, Example II, III],

$$A_6 = \begin{bmatrix} -3.5i & -12 & 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & -2.5i & -8 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & -1.5i & -4 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & -0.5i & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0.5i & 4 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1.5i & 8 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 2.5i & 12 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 3.5i \end{bmatrix},$$

$$A_7 = \begin{bmatrix} -3.5 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -2.5 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1.5 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -0.5 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.5 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1.5 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 2.5 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 3 \end{bmatrix}.$$

This matrix comes from [9, Test 2],

$$A_8 = \begin{bmatrix} -149 & -50 & -154 \\ 537 & 180 & 546 \\ -27 & -9 & -25 \end{bmatrix}.$$

The next matrix is taken from [14, Example 6.3],

$$A_9 = \begin{bmatrix} 1 + 10^{-7} & 10^5 & 10^4 \\ 0 & 1 + 10^{-1} & 10^5 \\ 0 & 0 & 11 \end{bmatrix}.$$

Additionally, the following four matrices were provided to us by Awad Al-Mohy from his algorithm test set,

$$A_{10} = \begin{bmatrix} \exp(i(\pi - 10^{-7})) & 1 \\ 0 & \exp(i(\pi + 10^{-7})) \end{bmatrix}, \quad A_{11} = \begin{bmatrix} \exp(i(\pi - 10^{-7})) & 1000 \\ 0 & \exp(i(\pi + 10^{-7})) \end{bmatrix},$$

$$A_{12} = \begin{bmatrix} \exp(i(\pi - 10^{-7})) & 1 \\ 0 & (1 + 10^{-7}i) \cdot \exp(i(\pi - 10^{-7})) \end{bmatrix},$$
$$A_{13} = \begin{bmatrix} \exp(i(\pi - 10^{-7})) & 1000 \\ 0 & (1 + 10^{-7}i) \cdot \exp(i(\pi - 10^{-7})) \end{bmatrix}.$$

REFERENCES

- [1] A.H. Al-Mohy. Conditioning of matrix functions of quasi-triangular matrices. *Technical Report*, 2022.
- [2] A.H. Al-Mohy and B. Arslan. The complex step approximation to the higher order Fréchet derivatives of a matrix function. *Numer. Algorithms*, 87(3):1061–1074, 2021.
- [3] A.H. Al-Mohy and N.J. Higham. A new scaling and squaring algorithm for the matrix exponential. *SIAM J. Matrix Anal. Appl.*, 31(3):970–989, 2009.
- [4] A.H. Al-Mohy and N.J. Higham. Computing the action of the matrix exponential, with an application to exponential integrators. *SIAM J. Sci. Comput.*, 33(2):488–511, 2011.
- [5] A.H. Al-Mohy and N.J. Higham. Improved inverse scaling and squaring algorithms for the matrix logarithm. *SIAM J. Sci. Comput.*, 34(4):C153–C169, 2012.
- [6] V. Arsigny, O. Commowick, and N. Ayache. A fast and log-Euclidean polyaffine framework for locally linear registration. *J. Math. Imag. Vis.*, 33:222–238, 2009.
- [7] B. Arslan, V. Noferini, and F. Tisseur. The structured condition number of a differentiable map between matrix manifolds, with applications. *SIAM J. Matrix Anal. Appl.*, 40(2):774–799, 2019.
- [8] M. Benzi, E. Estrada, and C. Klymko. Ranking hubs and authorities using matrix functions. *Linear Algebra Appl.*, 438(5):2447–2474, 2013.
- [9] J.R. Cardoso and F.S. Leite. Theoretical and numerical considerations about Padé approximants for the matrix logarithm. *Linear Algebra Appl.*, 330(1-3):31–42, 2001.
- [10] P.I. Davies. Structured conditioning of matrix functions. *Electron. J. Linear Algebra*, 11:132–161, 2004.
- [11] P.I. Davies and N.J. Higham. A Schur-Parlett algorithm for computing matrix functions. *SIAM J. Matrix Anal. Appl.*, 25(2):464–485, 2003.
- [12] E. Deadman, N.J. Higham, and R. Ralha. A recursive blocked Schur algorithm for computing the matrix square root. In *Proceedings of the 11th International Conference on Applied Parallel and Scientific Computing*, 171–182, 2012.
- [13] J.W. Demmel. On condition numbers and the distance to the nearest ill-posed problem. *Numer. Math.*, 51(3):251–289, 1987.
- [14] L. Dieci, B. Morini, and A. Papini. Computational techniques for real logarithms of matrices. *SIAM J. Matrix Anal. Appl.*, 17(3):570–593, 1996.
- [15] L. Dieci and A. Papini. Padé approximation for the exponential of a block triangular matrix. *Linear Algebra Appl.*, 308(1–3):183–202, 2000.
- [16] P. Grindrod and D.J. Higham. A dynamical systems view of network centrality. *Proc. Roy. Soc. London Ser.*, 470(2165), 2014.
- [17] D.J. Higham. Condition numbers and their condition numbers. *Linear Algebra Appl.*, 214(0):193 – 213, 1995.
- [18] N.J. Higham. *The Matrix Computation Toolbox*. <http://www.ma.man.ac.uk/~higham/mctoolbox>.
- [19] N.J. Higham. *Functions of Matrices: Theory and Computation*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2008.
- [20] N.J. Higham and L. Lin. On p th roots of stochastic matrices. *Linear Algebra Appl.*, 435(3):448–463, 2011.
- [21] N.J. Higham and X. Liu. A multiprecision derivative-free Schur-Parlett algorithm for computing matrix functions. *SIAM J. Matrix Anal. Appl.*, 42(3):1401–1422, 2021.
- [22] N.J. Higham and S.D. Relton. Higher order Fréchet derivatives of matrix functions and the level-2 condition number. *SIAM J. Matrix Anal. Appl.*, 35(3):1019–1037, 2014.
- [23] R.B. Israel, J.S. Rosenthal, and J.Z. Wei. Finding generators for Markov chains via empirical transition matrices, with applications to credit ratings. *Math. Finance*, 11(2):245–265, 2001.
- [24] D.P. Jagger. MATLAB toolbox for classical matrix groups. M.Sc. Thesis, University of Manchester, Manchester, England, 2003.



- [25] C. Kenney and A.J. Laub. Condition estimates for matrix functions. *SIAM J. Matrix Anal. Appl.*, 10(2):191–209, 1989.
- [26] J.C. Lagarias, J.A. Reeds, M.H. Wright, and P.E. Wright. Convergence properties of the Nelder–Mead simplex method in low dimensions. *SIAM J. Optim.*, 9(1):112–147, 1998.
- [27] M. Michel and S. Zell. TimeEvolver: A program for time evolution with improved error bound. *Comput. Phys. Commun.*, 277:108374, 2022.
- [28] S.T. Parker, D.M. Lorenzetti, and M.D. Sohn. Implementing state-space methods for multizone contaminant transport. *Buil. Environ.*, 71:131–139, 2014.
- [29] B.N. Parlett and K.C. Ng. Development of an accurate algorithm for $\exp(Bt)$. Technical report, California Univ. Berkeley Center for Pure and Applied Mathematics, 1985.
- [30] M. Schweitzer. Integral representations for higher-order Fréchet derivatives of matrix functions: Quadrature algorithms and new results on the level-2 condition number. *Linear Algebra Appl.*, 656:247–276, 2023.
- [31] F. Tisseur and S. Graillat. Structured condition numbers and backward errors in scalar product spaces. *Electron. J. Linear Algebra*, 15:159–177, 2006.
- [32] C. Van Loan. The sensitivity of the matrix exponential. *SIAM J. Numer. Anal.*, 14(6):971–981, 1977.
- [33] R.C. Ward. Numerical computation of the matrix exponential with accuracy estimate. *SIAM J. Numer. Anal.*, 14(4):600–610, 1977.