Pseudospectra of Loewner Matrix Pencils

Mark Embree and A. Cosmin Ioniţă

Abstract Loewner matrix pencils play a central role in the system realization theory of Mayo and Antoulas, an important development in data-driven modeling. The eigenvalues of these pencils reveal system poles. How robust are the poles recovered via Loewner realization? With several simple examples, we show how pseudospectra of Loewner pencils can be used to investigate the influence of interpolation point location and partitioning on pole stability, the transient behavior of the realized system, and the effect of noisy measurement data. We include an algorithm to efficiently compute such pseudospectra by exploiting Loewner structure.

Keywords Generalized eigenvalue problem · Pseudospectra · System realization

1 Introduction

The landmark systems realization theory of Mayo and Antoulas [19] shows how to construct a dynamical system that interpolates tangential frequency domain measurements of a multi-input, multi-output system. Central to this development is the *matrix pencil* $z\mathbb{L} - \mathbb{L}_s$ composed of Loewner and shifted Loewner matrices $\mathbb{L}$ and $\mathbb{L}_s$ that encode the interpolation data. When this technique is used for exact system recovery (as opposed to data-driven model reduction), the eigenvalues of this pencil match the poles of the transfer function of the original system. However other spectral properties, including the sensitivity of the eigenvalues to perturbation, can differ greatly, depending on the location of the interpolation points in the complex plane

---

Dedicated to Thanos Antoulas.

M. Embree (✉)
Department of Mathematics and Computational Modeling and Data Analytics Division, Academy of Integrated Science, Virginia Tech, Blacksburg, VA 24061, USA
e-mail: embree@vt.edu

A. C. Ioniţă
The MathWorks Inc., Natick, MA 01760, USA
e-mail: cionita@mathworks.com

relative to the system poles, and how the interpolation points are partitioned. Since one uses $z\mathbb{L} - \mathbb{L}_s$ to learn about the original system, such subtle differences matter.

Pseudospectra are sets in the complex plane that contain the eigenvalues but provide additional insight about the sensitivity of those eigenvalues to perturbation and the transient behavior of the underlying dynamical system. While most often used to analyze single matrices, pseudospectral concepts have been extended to matrix pencils (generalized eigenvalue problems).

This introductory note shows several ways to use pseudospectra to investigate spectral questions involving Loewner pencils derived from system realization problems. Using simple examples, we explore the following questions.

- How do the locations of the interpolation points and their partition into "left" and "right" points affect the sensitivity of the eigenvalues of $z\mathbb{L} - \mathbb{L}_s$?
- Do solutions to the dynamical system $\mathbb{L}\dot{\mathbf{x}}(t) = \mathbb{L}_s\mathbf{x}(t)$ mimic solutions to the original system $\mathbf{E}\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t)$, especially in the transient regime? Does this agreement depend on the interpolation points?
- How do noisy measurements affect the eigenvalues of $z\mathbb{L} - \mathbb{L}_s$?

We include an algorithm for computing pseudospectra of an n-dimensional Loewner pencil in $\mathcal{O}(n^2)$ operations, improving the $\mathcal{O}(n^3)$ cost for generic matrix pencils; the appendix gives a MATLAB implementation.

Throughout this note, we use $\sigma(\cdot)$ to denote the spectrum (eigenvalues) of a matrix or matrix pencil, and $\|\cdot\|$ to denote the vector 2-norm and the matrix norm it induces. (All definitions here can readily be adapted to other norms, as needed. The algorithm, however, is designed for use with the 2-norm.)

2 Loewner Realization Theory in a Nutshell

We briefly summarize Loewner realization theory, as developed by Mayo and Antoulas [19]; see also [1]. Consider the linear, time-invariant dynamical system

$$\mathbf{E}\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t), \qquad \mathbf{y}(t) = \mathbf{C}\mathbf{x}(t)$$

for $\mathbf{A}, \mathbf{E} \in \mathbb{C}^{n \times n}, \mathbf{B} \in \mathbb{C}^{n \times m}$, and $\mathbf{C} \in \mathbb{C}^{p \times n}$, with which we associate, via the Laplace transform, the transfer function $\mathbf{H}(z) = \mathbf{C}(z\mathbf{E} - \mathbf{A})^{-1}\mathbf{B}$.

Given *tangential measurements* of $\mathbf{H}(z)$ we seek a realization of the system that interpolates the given data. More precisely, consider the *right interpolation data*

- distinct interpolation points $\lambda_1, \ldots, \lambda_\varrho \in \mathbb{C}$;
- interpolation directions $\mathbf{r}_1, \ldots, \mathbf{r}_\varrho \in \mathbb{C}^m$;
- function values $\mathbf{w}_1, \ldots, \mathbf{w}_\varrho \in \mathbb{C}^p$;

and *left interpolation data*

- distinct interpolation points $\mu_1, \ldots, \mu_\nu \in \mathbb{C}$;

- interpolation directions $\boldsymbol{\ell}_1, \ldots, \boldsymbol{\ell}_\nu \in \mathbb{C}^p$;
- function values $\mathbf{v}_1, \ldots, \mathbf{v}_\nu \in \mathbb{C}^m$.

Assume the left and right interpolation points are disjoint, $\{\lambda_i\}_{i=1}^\varrho \cap \{\mu_j\}_{j=1}^\nu = \emptyset$; (in our examples, all $\varrho + \nu$ points are distinct).

The interpolation problem seeks matrices $\widehat{\mathbf{A}}, \widehat{\mathbf{E}}, \widehat{\mathbf{B}}, \widehat{\mathbf{C}}$ for which the transfer function $\widehat{\mathbf{H}}(z) = \widehat{\mathbf{C}}(z\widehat{\mathbf{E}} - \widehat{\mathbf{A}})^{-1}\widehat{\mathbf{B}}$ interpolates the data: for $i = 1, \ldots, \varrho$ and $j = 1, \ldots, \nu$,

$$\widehat{\mathbf{H}}(\lambda_i)\mathbf{r}_i = \mathbf{w}_i, \qquad \boldsymbol{\ell}_j^*\widehat{\mathbf{H}}(\mu_j) = \mathbf{v}_j^*.$$

Two structured matrices play a crucial role in the development of Mayo and Antoulas [19]. From the data, construct the *Loewner* and *shifted Loewner* matrices

$$\mathbb{L} = \begin{bmatrix} \frac{\mathbf{v}_1^*\mathbf{r}_1 - \boldsymbol{\ell}_1^*\mathbf{w}_1}{\mu_1 - \lambda_1} & \cdots & \frac{\mathbf{v}_1^*\mathbf{r}_\varrho - \boldsymbol{\ell}_1^*\mathbf{w}_\varrho}{\mu_1 - \lambda_\varrho} \\ \vdots & \ddots & \vdots \\ \frac{\mathbf{v}_\nu^*\mathbf{r}_1 - \boldsymbol{\ell}_\nu^*\mathbf{w}_1}{\mu_\nu - \lambda_1} & \cdots & \frac{\mathbf{v}_\nu^*\mathbf{r}_\varrho - \boldsymbol{\ell}_\nu^*\mathbf{w}_\varrho}{\mu_\nu - \lambda_\varrho} \end{bmatrix}, \quad \mathbb{L}_s = \begin{bmatrix} \frac{\mu_1\mathbf{v}_1^*\mathbf{r}_1 - \lambda_1\boldsymbol{\ell}_1^*\mathbf{w}_1}{\mu_1 - \lambda_1} & \cdots & \frac{\mu_1\mathbf{v}_1^*\mathbf{r}_\varrho - \lambda_\varrho\boldsymbol{\ell}_1^*\mathbf{w}_\varrho}{\mu_1 - \lambda_\varrho} \\ \vdots & \ddots & \vdots \\ \frac{\mu_\nu\mathbf{v}_\nu^*\mathbf{r}_1 - \lambda_1\boldsymbol{\ell}_\nu^*\mathbf{w}_1}{\mu_\nu - \lambda_1} & \cdots & \frac{\mu_\nu\mathbf{v}_\nu^*\mathbf{r}_\varrho - \lambda_\varrho\boldsymbol{\ell}_\nu^*\mathbf{w}_\varrho}{\mu_\nu - \lambda_\varrho} \end{bmatrix}, \quad (1)$$

i.e., the (i, j) entries of these $\nu \times \varrho$ matrices have the form

$$(\mathbb{L})_{i,j} = \frac{\mathbf{v}_i^*\mathbf{r}_j - \boldsymbol{\ell}_i^*\mathbf{w}_j}{\mu_i - \lambda_j}, \qquad (\mathbb{L}_s)_{i,j} = \frac{\mu_i\mathbf{v}_i^*\mathbf{r}_j - \lambda_j\boldsymbol{\ell}_i^*\mathbf{w}_j}{\mu_i - \lambda_j}.$$

Now collect the data into matrices. The right interpolation points, directions, and data are stored in

$$\boldsymbol{\Lambda} = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_\varrho \end{bmatrix} \in \mathbb{C}^{\varrho \times \varrho}, \quad \mathbf{R} = \begin{bmatrix} \mathbf{r}_1 \cdots \mathbf{r}_\varrho \end{bmatrix} \in \mathbb{C}^{m \times \varrho}, \quad \mathbf{W} = \begin{bmatrix} \mathbf{w}_1 \cdots \mathbf{w}_\varrho \end{bmatrix} \in \mathbb{C}^{p \times \varrho},$$

while the left interpolation points, directions, and data are stored in

$$\mathbf{M} = \begin{bmatrix} \mu_1 & & \\ & \ddots & \\ & & \mu_\nu \end{bmatrix} \in \mathbb{C}^{\nu \times \nu}, \quad \mathbf{L} = \begin{bmatrix} \boldsymbol{\ell}_1 \cdots \boldsymbol{\ell}_\nu \end{bmatrix} \in \mathbb{C}^{p \times \nu}, \quad \mathbf{V} = \begin{bmatrix} \mathbf{v}_1 \cdots \mathbf{v}_\nu \end{bmatrix} \in \mathbb{C}^{m \times \nu}.$$

2.1   Selecting and Arranging Interpolation Points

As Mayo and Antoulas observe, Sylvester equations connect these matrices:

$$\mathbb{L}\boldsymbol{\Lambda} - \mathbf{M}\mathbb{L} = \mathbf{L}^*\mathbf{W} - \mathbf{V}^*\mathbf{R}, \qquad \mathbb{L}_s\boldsymbol{\Lambda} - \mathbf{M}\mathbb{L}_s = \mathbf{L}^*\mathbf{W}\boldsymbol{\Lambda} - \mathbf{M}\mathbf{V}^*\mathbf{R}. \tag{2}$$

Just using the dimensions of the components, note that

$$\operatorname{rank}(\mathbf{L}^*\mathbf{W}), \; \operatorname{rank}(\mathbf{L}^*\mathbf{W}\boldsymbol{\Lambda}), \; \leq \min\{\nu, \varrho, p\}$$
$$\operatorname{rank}(\mathbf{V}^*\mathbf{R}), \; \operatorname{rank}(\mathbf{V}^*\mathbf{R}\mathbf{M}) \leq \min\{\nu, \varrho, m\}.$$

Thus for modest m and p, the Sylvester equations (2) must have low-rank right-hand sides.[1] This situation often implies the rapid decay of singular values of solutions to the Sylvester equation [2, 3, 21, 22, 26]. While this phenomenon is convenient for balanced truncation model reduction (enabling low-rank approximations to the controllability and observability Gramians), it is less welcome in the Loewner real-ization setting, where the rank of $\mathbb{L}$ should reveal the order of the original system: *fast decay of the singular values of $\mathbb{L}$ makes this rank ambiguous*. Since $\boldsymbol{\Lambda}$ and $\mathbf{M}$ are diagonal, they are normal matrices, and hence Theorem 2.1 of [3] gives

$$\frac{s_{qk+1}(\mathbb{L})}{s_1(\mathbb{L})} \leq \inf_{\phi \in \mathcal{R}_{k,k}} \frac{\max\{|\phi(\lambda)| : \lambda \in \{\lambda_1, \ldots, \lambda_\varrho\}\}}{\min\{|\phi(\mu)| : \mu \in \{\mu_1, \ldots, \mu_\nu\}\}}, \tag{3}$$

where $s_j(\cdot)$ denotes the jth largest singular value, $q = \operatorname{rank}(\mathbf{L}^*\mathbf{W} - \mathbf{V}^*\mathbf{R})$, and $\mathcal{R}_{k,k}$ denotes the set of irreducible rational functions whose numerators and denominators are polynomials of degree k or less.[2] The right hand side of (3) will be *small* when there exists some $\phi \in \mathcal{R}_{k,k}$ for which all $|\phi(\lambda_i)|$ are small, while all $|\phi(\mu_j)|$ are large: a good separation of $\{\lambda_i\}$ from $\{\mu_j\}$ is thus *sufficient* to ensure the rapid decay of the singular values of $\mathbb{L}$ and $\mathbb{L}_s$. (Beckermann and Townsend give an explicit bound for the singular values of Loewner matrices when the interpolation points fall in disjoint real intervals [3, Corollary 4.2].)

Here, if we want the singular values of $\mathbb{L}$ and $\mathbb{L}_s$ to reveal the system's order (without decay of singular values as an accident of the arrangement of interpolation points), it is *necessary* for one $|\phi(\lambda_i)|$ to be about the same size as the smallest value of $|\phi(\mu_j)|$ for *all* $\phi \in \mathcal{R}_{k,k}$. Roughly speaking, we want the left and right interpolation points to be close together (even interleaved). While this arrangement is *necessary* for slow decay of the singular values, it does not alone prevent such decay, as will be seen in examples in Fig. 4 (since (3) is only an upper bound).

Another heuristic, based on the Cauchy-like structure of $\mathbb{L}$ and $\mathbb{L}_s$, also suggests the left and right interpolation points should be close together. Namely, $\mathbb{L}$ and $\mathbb{L}_s$ are

---

[1] For single-input, single-output (SISO) systems, $m = p = 1$, so the rank of the right-hand sides in (2) cannot exceed two. The same will apply for multi-input, multi-output systems with identical left and right interpolation directions: $\boldsymbol{\ell}_i \equiv \boldsymbol{\ell}$ for all $i = 1, \ldots, \nu$ and $\mathbf{r}_j \equiv \mathbf{r}$ for all $j = 1, \ldots, \varrho$.
[2] The same bound holds for $s_{qk+1}(\mathbb{L}_s)/s_1(\mathbb{L}_s)$ with $q = \operatorname{rank}(\mathbf{L}^*\mathbf{W}\boldsymbol{\Lambda} - \mathbf{M}\mathbf{V}^*\mathbf{R})$.

a more general form of the Cauchy matrix $(\mathcal{C})_{i,j} = 1/(\mu_i - \lambda_j)$, whose determinant has the elegant formula (e.g., [16, p. 38])

$$\det(\mathcal{C}) = \frac{\prod_{1 \le i < j \le n}(\mu_j - \mu_i)(\lambda_i - \lambda_j)}{\prod_{1 \le i \le j \le n}(\mu_i - \lambda_j)}. \tag{4}$$

It is an open question if $\det(\mathbb{L})$ and $\det(\mathbb{L}_s)$ have similarly elegant formulas. Nevertheless, $\det(\mathbb{L})$ and $\det(\mathbb{L}_s)$ do have the same denominator as $\det(\mathcal{C})$ (which can be checked by recursively subtracting the first row from all other rows when computing the determinant). This observation suggests that to avoid artificially small determinants for $\mathbb{L}$ and $\mathbb{L}_s$ (which, up to sign, are the products of the singular values) it is *necessary* for the denominator of (4) to be small, and, thus, for the left and right interpolation points to be close together.

In practice, we often start with initial interpolation points $x_1, \dots, x_{2n}$ that we want to partition into left and right interpolation points to form $\mathbb{L}$ and $\mathbb{L}_s$. Our analysis of (4) suggests a simple way to arrange the interpolation points such that the denominator of $\det(\mathbb{L})$ and $\det(\mathbb{L}_s)$ is small: relabel the points to satisfy

$$x_k = \underset{k \le q \le 2n}{\arg\min}\, |x_{k-1} - x_q|, \quad \text{for } k = 2, \dots, 2n, \tag{5}$$

$$\{\mu_i\} = \{x_1, x_3, \dots, x_{2n-1}\}, \quad \{\lambda_j\} = \{x_2, x_4, \dots, x_{2n}\}.$$

The greedy reordering in (5) ensures that $|x_k - x_{k-1}|$ is small and allows us to simply interleave the left and right interpolation points. Moreover, when $x_1, \dots, x_{2n}$ are located on a line, the reordering in (5) simplifies to directly interleaving μ_i and λ_j and, thus, it can be skipped. This ordering need not be optimal, as we do not visit all possible combinations of $\mu_i - \lambda_j$; it simply seeks a partition that yields a large determinant (which must also depend on the interpolation *data*). We note its simplicity, effectiveness, and efficiency (requiring only $\mathcal{O}(n^2)$ operations). (For SISO systems, the state-space representation is equivalent to a barycentric form [17, p. 77]. We thank the referee for observing that in barycentric rational Remez approximations [9], a similar interleaving of support points from the set of reference points produces favorable numerical conditioning. Recent work by Gosea and Antoulas gives further numerical evidence for interleaving Loewner points [13].)

2.2 Construction of Interpolants

Throughout we make the fundamental assumptions that for all $\widehat{z} \in \{\lambda_i\}_{i=1}^{\varrho} \cup \{\mu_j\}_{j=1}^{\nu}$,

$$\text{rank}(\widehat{z}\mathbb{L} - \mathbb{L}_s) = \text{rank}\left(\begin{bmatrix} \mathbb{L} \\ \mathbb{L}_s \end{bmatrix}\right) = \text{rank}\left(\begin{bmatrix} \mathbb{L} & \mathbb{L}_s \end{bmatrix}\right) =: r,$$

and we presume the underlying dynamical system is controllable and observable.

When $r = \nu = \varrho$ is the order of the system, Mayo and Antoulas [19, Lemma 5.1] show that the transfer function $\widehat{\mathbf{H}}(z) := \widehat{\mathbf{C}}(z\widehat{\mathbf{E}} - \widehat{\mathbf{A}})^{-1}\widehat{\mathbf{B}}$ defined by

$$\widehat{\mathbf{E}} = -\mathbb{L}, \quad \widehat{\mathbf{A}} = -\mathbb{L}_s, \quad \widehat{\mathbf{B}} = \mathbf{V}^*, \quad \widehat{\mathbf{C}} = \mathbf{W} \tag{6}$$

interpolates the $\varrho + \nu$ data values.

When $r < \max(\nu, \varrho)$, fix some $\widehat{z} \in \{\lambda_i\}_{i=1}^{\varrho} \cup \{\mu_j\}_{j=1}^{\nu}$ and compute the (economy-sized) singular value decomposition

$$\widehat{z}\mathbb{L} - \mathbb{L}_s = \mathbf{Y}\mathbf{\Sigma}\mathbf{X}^*,$$

with $\mathbf{Y} \in \mathbb{C}^{\nu \times r}$, $\mathbf{\Sigma} \in \mathbb{R}^{r \times r}$, and $\mathbf{X} \in \mathbb{C}^{\varrho \times r}$. Then with

$$\widehat{\mathbf{E}} = -\mathbf{Y}^*\mathbb{L}\mathbf{X}, \quad \widehat{\mathbf{A}} = -\mathbf{Y}^*\mathbb{L}_s\mathbf{X}, \quad \widehat{\mathbf{B}} = \mathbf{Y}^*\mathbf{V}^*, \quad \widehat{\mathbf{C}} = \mathbf{W}\mathbf{X}, \tag{7}$$

$\widehat{\mathbf{H}}(z) := \widehat{\mathbf{C}}(z\widehat{\mathbf{E}} - \widehat{\mathbf{A}})^{-1}\widehat{\mathbf{B}}$ gives an order-r system interpolating the data [19, Theorem 5.1].

3 Pseudospectra for Matrix Pencils

Though introduced decades earlier, in the 1990s pseudospectra emerged as a popular tool for analyzing the behavior of dynamical systems (see, e.g., [30]), eigenvalue perturbations (see, e.g., [4]), and stability of uncertain linear time-invariant (LTI) systems (see, e.g., [15]).

Definition 1 For a matrix $\mathbf{A} \in \mathbb{C}^{n \times n}$ and $\varepsilon > 0$, the *ε-pseudospectrum* of $\mathbf{A}$ is

$$\sigma_\varepsilon(\mathbf{A}) = \{z \in \mathbb{C} \text{ is an eigenvalue of } \mathbf{A} + \mathbf{\Gamma} \text{ for some } \mathbf{\Gamma} \in \mathbb{C}^{n \times n} \text{ with } \|\mathbf{\Gamma}\| < \varepsilon\}. \tag{8}$$

For all $\varepsilon > 0$, $\sigma_\varepsilon(\mathbf{A})$ is a bounded, open subset of the complex plane that contains the eigenvalues of $\mathbf{A}$. (A popular variation uses the weak inequality $\|\mathbf{\Gamma}\| \leq \varepsilon$; the strict inequality has favorable properties for operators on infinite-dimensional spaces [5].) Definition 1 motivates pseudospectra via eigenvalues of perturbed matrices. A numerical analyst studying accuracy of a backward stable eigenvalue algorithm might be concerned with ε on the order of $n\|\mathbf{A}\|\varepsilon_{\text{mach}}$, where $\varepsilon_{\text{mach}}$ denotes the *machine epsilon* for the floating point system [20]. An engineer or scientist might consider $\sigma_\varepsilon(\mathbf{A})$ for much larger ε values, corresponding to uncertainty in parameters or data that contribute to the entries of $\mathbf{A}$.

Via the singular value decomposition, one can show that (8) is equivalent to

$$\sigma_\varepsilon(\mathbf{A}) = \{z \in \mathbb{C} : \|(z\mathbf{I} - \mathbf{A})^{-1}\| > 1/\varepsilon\}; \tag{9}$$

see, e.g., [29, Chap. 2]. The presence of the resolvent $(z\mathbf{I} - \mathbf{A})^{-1}$ in this definition suggests a connection to the transfer function $\mathbf{H}(z) = \mathbf{C}(z\mathbf{I} - \mathbf{A})^{-1}\mathbf{B}$ for the system

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t), \qquad \mathbf{y}(t) = \mathbf{C}\mathbf{x}(t).$$

Indeed, definition (1) readily leads to bounds on $\|e^{t\mathbf{A}}\|$, and hence transient growth of solutions to $\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t)$; see [29, Part IV].

Various extensions of pseudospectra have been proposed to handle more general eigenvalue problems and dynamical systems; see [7] for a concise survey. The first elaborations addressed the generalized eigenvalue problem $\mathbf{A}\mathbf{x} = \lambda\mathbf{E}\mathbf{x}$ (i.e., the matrix pencil $z\mathbf{E} - \mathbf{A}$) [10, 23, 24]. Here we focus on the definition proposed by Fraysségé, Gueury, Nicoud, and Toumazou [10], which is ideally suited to analyzing eigenvalues of nearby matrix pencils. To permit the perturbations to $\mathbf{A}$ and $\mathbf{E}$ to be scaled independently, this definition includes two additional parameters, γ and δ.

Definition 2 Let $\gamma, \delta > 0$. For a pair of matrices $\mathbf{A}, \mathbf{E} \in \mathbb{C}^{n \times n}$ and any $\varepsilon > 0$, the ε-(γ, δ)-*pseudospectrum* $\sigma_\varepsilon^{(\gamma,\delta)}(\mathbf{A}, \mathbf{E})$ of the matrix pencil $z\mathbf{E} - \mathbf{A}$ is the set

$$\sigma_\varepsilon^{(\gamma,\delta)}(\mathbf{A}, \mathbf{E}) = \{z \in \mathbb{C} \text{ is an eigenvalue of the pencil } z(\mathbf{E} + \mathbf{\Delta}) - (\mathbf{A} + \mathbf{\Gamma})$$
$$\text{for some } \mathbf{\Gamma}, \mathbf{\Delta} \in \mathbb{C}^{n \times n} \text{ with } \|\mathbf{\Gamma}\| < \varepsilon\gamma, \|\mathbf{\Delta}\| < \varepsilon\delta\}.$$

This definition has been extended to matrix polynomials in [14, 27].

Remark 1 Note that $\sigma_\varepsilon^{(\gamma,\delta)}(\mathbf{A}, \mathbf{E})$ is an open, nonempty subset of the complex plane, but it need not be bounded.

(a) If $z\mathbf{E} - \mathbf{A}$ is a singular pencil (rank$(z\mathbf{E} - \mathbf{A}) < n$ for all $z \in \mathbb{C}$), then $\sigma(\mathbf{A}, \mathbf{E}) = \mathbb{C}$.
(b) If $z\mathbf{E} - \mathbf{A}$ is a regular pencil but $\mathbf{E}$ is not invertible, then $\sigma(\mathbf{A}, \mathbf{E})$ contains an infinite eigenvalue, and hence $\sigma_\varepsilon^{(\gamma,\delta)}(\mathbf{A}, \mathbf{E})$ is unbounded.
(c) If $\mathbf{E}$ is nonsingular but $\varepsilon\delta$ exceeds the distance of $\mathbf{E}$ to singularity (the smallest singular value of $\mathbf{E}$), then $\sigma_\varepsilon^{(\gamma,\delta)}(\mathbf{A}, \mathbf{E})$ contains the point at infinity.

Since these pseudospectra can be unbounded, Lavallée [18] and Higham and Tisseur [14] visualize $\sigma_\varepsilon^{(\gamma,\delta)}(\mathbf{A}, \mathbf{E})$ as stereographic projections on the Riemann sphere.

Remark 2 Just as the conventional pseudospectrum $\sigma_\varepsilon(\mathbf{A})$ can be characterized using the resolvent of $\mathbf{A}$ in (9), Frayseé et al. [10] show that Definition 2 is equivalent to

$$\sigma_\varepsilon^{(\gamma,\delta)}(\mathbf{A}, \mathbf{E}) = \left\{z \in \mathbb{C} : \|(z\mathbf{E} - \mathbf{A})^{-1}\| > \frac{1}{\varepsilon(\gamma + |z|\delta)}\right\}. \qquad (10)$$

This formula suggests a way to compute $\sigma_\varepsilon^{(\gamma,\delta)}(\mathbf{A}, \mathbf{E})$: evaluate $\|(z\mathbf{E} - \mathbf{A})^{-1}\|$ on a grid of points covering a relevant region of the complex plane (or the Riemann sphere) and use a contour plotting routine to draw boundaries of $\sigma_\varepsilon^{(\gamma,\delta)}(\mathbf{A}, \mathbf{E})$. The accuracy of the resulting pseudospectra depends on the density of the grid. Expedient algorithms

for computing $\|(z\mathbf{E} - \mathbf{A})^{-1}\|$ can be derived by computing a unitary simultaneous triangularization (generalized Schur form) of $\mathbf{A}$ and $\mathbf{E}$ in $\mathcal{O}(n^3)$ operations, then using inverse iteration or inverse Lanczos, as described by Trefethen [28] and Wright [31], to compute $\|(z\mathbf{E} - \mathbf{A})^{-1}\|$ at each grid point in $\mathcal{O}(n^2)$ operations. For the structured Loewner pencils of interest here, one can compute $\|(z\mathbf{E} - \mathbf{A})^{-1}\|$ in $\mathcal{O}(n^2)$ operations without recourse to the $\mathcal{O}(n^3)$ preprocessing step, as proposed in Sect. 4.

Remark 3 Definition 2 can be extended to $\delta = 0$ by only perturbing $\mathbf{A}$ [23]:

$$
\begin{aligned}
\sigma_\varepsilon^{(1,0)}(\mathbf{A}, \mathbf{E}) = \{&z \in \mathbb{C} \text{ is an eigenvalue of the pencil } z\mathbf{E} - (\mathbf{A} + \mathbf{\Gamma}) \\
&\text{for some } \mathbf{\Gamma} \in \mathbb{C}^{n \times n} \text{ with } \|\mathbf{\Gamma}\| < \varepsilon\}. \\
= \{&z \in \mathbb{C} : \|(z\mathbf{E} - \mathbf{A})^{-1}\| > 1/\varepsilon\}.
\end{aligned}
$$

This definition may be more suitable for cases where $\mathbf{E}$ is fixed and uncertainty in the system only emerges, e.g., through physical parameters that appear in $\mathbf{A}$.

Remark 4 Since we ultimately intend to study the pseudospectra $\sigma_\varepsilon^{(\gamma,\delta)}(\mathbb{L}_s, \mathbb{L})$ of Loewner matrix pencils, one might question the use of generic perturbations $\mathbf{\Gamma}, \mathbf{\Delta} \in \mathbb{C}^{n \times n}$ in Definition 2. Should we restrict $\mathbf{\Gamma}$ and $\mathbf{\Delta}$ to maintain Loewner structure, i.e., so that $\mathbb{L}_s + \mathbf{\Gamma}$ and $\mathbb{L} + \mathbf{\Delta}$ maintain the coupled shifted Loewner–Loewner form in (1)? Such sets are called *structured pseudospectra*.

Three considerations motivate the study of generic perturbations $\mathbf{\Gamma}, \mathbf{\Delta} \in \mathbb{C}^{n \times n}$: one practical, one speculative, and one philosophical. (a) Beyond repeatedly computing the eigenvalues of Loewner pencils with randomly perturbed data, no systematic method is known to compute the coupled Loewner structured pseudospectra, i.e., no analogue of the resolvent-like definition (10) is known. (b) Rump [25] showed that in many cases, preserving structure has little effect on the standard matrix pseudospectra. For example, the structured ε-pseudospectrum of a Hankel matrix $\mathbf{H}$ allowing only complex Hankel perturbations exactly matches the unstructured ε-pseudospectrum $\sigma_\varepsilon(\mathbf{H})$ based on generic complex perturbations [25, Theorem 4.3]. *Whether a similar results holds for Loewner structured pencils is an interesting open question.* (c) If one seeks to analyze the *behavior* of dynamical systems (as opposed to eigenvalues of nearby matrices), then generic perturbations give much greater insight; see [29, p. 456] for an example where *real-valued* perturbations do not move the eigenvalues much toward the imaginary axis (hence the real structured pseudospectra are benign), yet the stable system still exhibits strong transient growth.

As we shall see in Sect. 6, Definition 2 provides a helpful tool for investigating the sensitivity of eigenvalues of matrix pencils. A different generalization of Definition 1 gives insight into the transient behavior of solutions of $\mathbf{E}\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t)$. This approach is discussed in [29, Chap. 45], following [23, 24], and has been extended to handle singular $\mathbf{E}$ in [8] (for differential-algebraic equations and descriptor systems). Restricting our attention here to nonsingular $\mathbf{E}$, we analyze the conventional (single matrix) pseudospectra $\sigma_\varepsilon(\mathbf{E}^{-1}\mathbf{A})$. From these sets one can develop various upper and lower bounds on $\|e^{t\mathbf{E}^{-1}\mathbf{A}}\|$ and $\|\mathbf{x}(t)\|$ [29, Chap. 15]. Here we shall just

state one basic result. If $\sup\{\mathrm{Re}(z) : z \in \sigma_\varepsilon(\mathbf{E}^{-1}\mathbf{A})\} = K\varepsilon$ for some $K \geq 1$ (where $\mathrm{Re}(\cdot)$ denotes the real part of a complex number), then

$$\sup_{t \geq 0} \|e^{t\mathbf{E}^{-1}\mathbf{A}}\| \geq K. \tag{11}$$

This statement implies that there exists some unit-length initial condition $\mathbf{x}(0)$ such that $\|\mathbf{x}(t)\| \geq K$, even though $\sigma(\mathbf{A}, \mathbf{E})$ may be contained in the left half-plane. (Optimizing this bound over $\varepsilon > 0$ yields the Kreiss Matrix Theorem [29, (15.9)].)

Pseudospectra of matrix pencils provide a natural vehicle to explore that stability of the matrix pencil associated with the Loewner realization in (6). We shall thus investigate eigenvalue perturbations via $\sigma_\varepsilon^{(\gamma,\delta)}(\widehat{\mathbf{A}}, \widehat{\mathbf{E}}) = \sigma_\varepsilon^{(\gamma,\delta)}(\mathbb{L}_s, \mathbb{L})$ and transient behavior via $\sigma_\varepsilon(\mathbb{L}^{-1}\mathbb{L}_s)$.

4 Efficient Computation of Loewner Pseudospectra

We first present a novel technique for efficiently computing pseudospectra of large Loewner matrix pencils, $\sigma_\varepsilon^{(\gamma,\delta)}(\mathbb{L}_s, \mathbb{L})$, using the equivalent definition given in (10). When the Loewner matrix $\mathbb{L}$ is nonsingular, we employ inverse iteration to exploit the structure of the Loewner pencil to compute $\|(z\mathbb{L} - \mathbb{L}_s)^{-1}\|$ (in the two-norm) using only $\mathcal{O}(n^2)$ operations. This avoids the need to compute an initial simultaneous unitary triangularization of $\mathbb{L}_s$ and $\mathbb{L}$ using the QZ algorithm, an $\mathcal{O}(n^3)$ operation.

Inverse iteration (and inverse Lanczos) for $\|(z\mathbb{L} - \mathbb{L}_s)^{-1}\|$ requires computing

$$(z\mathbb{L} - \mathbb{L}_s)^{-*}(z\mathbb{L} - \mathbb{L}_s)^{-1}\mathbf{u} \tag{12}$$

for a series of vectors $\mathbf{u} \in \mathbb{C}^n$ (e.g., see [29, Chap. 39]). We invoke a property observed by Mayo and Antoulas [19], related to (2): by construction, the Loewner and shifted Loewner matrices satisfy $\mathbb{L}_s - \mathbb{L}\mathbf{\Lambda} = \mathbf{V}^*\mathbf{R}$. Thus the resolvent can be expressed using only $\mathbb{L}$ and not $\mathbb{L}_s$:

$$(z\mathbb{L} - \mathbb{L}_s)^{-1} = (\mathbb{L}(z\mathbf{I} - \mathbf{\Lambda}) - \mathbf{V}^*\mathbf{R})^{-1}.$$

We now use the Sherman–Morrison–Woodbury formula (see, e.g., [12]) to get

$$(z\mathbb{L} - \mathbb{L}_s)^{-1} = \left(\mathbf{I} + \mathbf{\Upsilon}(z)(\mathbf{I} - \mathbf{R}\mathbf{\Upsilon}(z))^{-1}\mathbf{R}\right)(z\mathbf{I} - \mathbf{\Lambda})^{-1}\mathbb{L}^{-1} = \mathbf{\Theta}(z)\mathbb{L}^{-1},$$

where $\mathbf{\Upsilon}(z) := (z\mathbf{I} - \mathbf{\Lambda})^{-1}\mathbb{L}^{-1}\mathbf{V}^*$ and $\mathbf{\Theta}(z) := \left(\mathbf{I} + \mathbf{\Upsilon}(z)(\mathbf{I} - \mathbf{R}\mathbf{\Upsilon}(z))^{-1}\mathbf{R}\right)(z\mathbf{I} - \mathbf{\Lambda})^{-1}$. As a result, we can compute the inverse iteration vectors in (12) as

$$(z\mathbb{L} - \mathbb{L}_s)^{-*}(z\mathbb{L} - \mathbb{L}_s)^{-1}\mathbf{u} = \mathbb{L}^{-*}\mathbf{\Theta}(z)^*\mathbf{\Theta}(z)\mathbb{L}^{-1}\mathbf{u}, \tag{13}$$

which requires solving several linear systems given by the same Loewner matrix $\mathbb{L}$, e.g., $\mathbb{L}^{-1}\mathbf{u}$, $\mathbb{L}^{-1}\mathbf{V}^*$.

Crucially, solving a linear system involving a Loewner matrix $\mathbb{L} \in \mathbb{C}^{n \times n}$ can be done efficiently in only $2(m + p + 1)n^2$ operations, since $\mathbb{L}$ has displacement rank $m + p$. More precisely, $\mathbb{L}$ is a Cauchy-like matrix that satisfies the Sylvester equation (2) given by diagonal generator matrices $\mathbf{\Lambda}$ and $\mathbf{M}$ and a right-hand side of rank at most $m + p$, i.e., $\mathrm{rank}(\mathbf{L}^*\mathbf{W} - \mathbf{V}^*\mathbf{R}) \leq m + p$. The displacement rank structure of $\mathbb{L}$ can be exploited to compute its LU factorization in only $2(m + p)n^2$ flops (see [11] and [12, Sect. 12.1]). Given the LU factorization of $\mathbb{L}$, solving $\mathbb{L}^{-1}\mathbf{u}$ in (13) via standard forward and backward substitution requires another $2n^2$ operations.

Next, multiplying $\mathbf{\Theta}(z)$ with the solution of $\mathbb{L}^{-1}\mathbf{u}$ requires a total of $2mn^2 + (4m^2 + 6m + 2)n + \frac{2}{3}m^3 - m^2$ operations, namely (to leading order on the factorizations):

- $2mn^2 + 2mn$ operations to compute $\mathbf{\Upsilon}(z) \in \mathbb{C}^{n \times m}$;
- $m^2(2n - 1) + m$ operations to compute $\mathbf{I} - \mathbf{R}\mathbf{\Upsilon}(z) \in \mathbb{C}^{m \times m}$;
- $\frac{2}{3}m^3 + 2m^2n$ operations to solve $(\mathbf{I} - \mathbf{R}\mathbf{\Upsilon}(z))^{-1}\mathbf{R} \in \mathbb{C}^{m \times n}$ via an LU factorization followed by n forward and backward substitutions;
- $n + m(2n - 1) + (2m - 1)n + 2n$ operations to multiply $\mathbf{\Theta}(z)$ with $\mathbb{L}^{-1}\mathbf{u}$.

Finally, multiplying with $\mathbb{L}^{-*}\mathbf{\Theta}(z)^*$ in (13) requires an additional $2n^2 + 2mn^2 + (4m^2 + 6m + 2)n + \frac{2}{3}m^3 - m^2$ operations, bringing the total cost of computing (13) to $2(3m + p + 1)n^2 + 4(2m^2 + 3m + 1)n + \frac{4}{3}m^3 - 2m^2$ operations.

In practice, the sizes of the right and left tangential directions are much smaller than the size of the Loewner pencil, i.e., $m, p \ll n$. For example, for scalar data (associated with SISO systems), $m = p = 1$. Therefore, in practice, computing (13) can be done in only $\mathcal{O}(n^2)$ operations.

Partial pivoting can be included in the LU factorization of the Loewner matrix $\mathbb{L}$ to overcome numerical difficulties. Adding partial pivoting maintains the $\mathcal{O}(n^2)$ operation count for the LU factorization of $\mathbb{L}$ (see [12, Sect. 12.1]), and hence computing (13) can still be done in $\mathcal{O}(n^2)$ operations. The appendix gives a MATLAB implementation of this efficient inverse iteration.

We measure these performance gains for a Loewner pencil generated by sampling $f(x) = \sum_{k=1}^{8} (-1)^{k+1} \left(1 + 100(x - k)^2 \right)^{-1/2} + (-1)^{k+1} \left(1 + 100(x - k - 1/2)^2 \right)^{-1/2}$ at $2n$ points uniformly spaced in the interval $[1, 8]$. We compare our new $\mathcal{O}(n^2)$ Loewner pencil inverse iteration against a standard implementation (see [29, p. 373]) applied to a simultaneous triangularization of $\mathbb{L}_s$ and $\mathbb{L}$. (The simultaneous triangularization costs $\mathcal{O}(n^3)$ but is fast, as MATLAB's qz routine invokes LAPACK code. For a fair comparison, we test against a C++ implementation of the fast Loewner code, compiled into a MATLAB $\mathrm{.mex}$ file.) Table 1 shows timings for both implementations by computing $\|(z\mathbb{L} - \mathbb{L}_s)^{-1}\|$ on a 200×200 grid of points. Exploiting the Loewner structure gives a significant performance improvement for large n.

We next examine two simple examples involving full-rank realization of SISO systems, to illustrate the kinds of insights one can draw from pseudospectra of Loewner pencils. (For these small examples we use the standard $\mathcal{O}(n^3)$ algorithm.)

Table 1 Comparison of speed of computing $\sigma_\varepsilon^{(1,1)}(\mathbb{L}_s, \mathbb{L})$ using the fast Loewner algorithm versus a generic inverse iteration method applied to simultaneous triangularizations of $\mathbb{L}_s$ and $\mathbb{L}$

n	$\mathcal{O}(n^2)$ algorithm (s)	$\mathcal{O}(n^3)$ algorithm (s)
100	1.85	1.65
200	6.75	6.75
300	17.24	33.30
400	49.15	65.05

5 Example 1: Eigenvalue Sensitivity and Transient Behavior

We first consider a simple controllable and observable SISO system with $n = 2$:

$$\mathbf{E} = \mathbf{I}, \quad \mathbf{A} = \begin{bmatrix} -1.1 & 1 \\ 1 & -1.1 \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \quad \mathbf{C} = \begin{bmatrix} 0 & 1 \end{bmatrix}. \quad (14)$$

This $\mathbf{A}$ is symmetric negative definite, with eigenvalues $\sigma(\mathbf{A}) = \{-0.1, -2.1\}$. Since the system is SISO, the transfer function $\mathbf{H}(s) = \mathbf{C}(s\mathbf{I} - \mathbf{A})^{-1}\mathbf{B}$ maps $\mathbb{C}$ to $\mathbb{C}$, and hence the choice of "interpolation directions" is trivial (though the division into "left" and "right" points matters). We take $\varrho = \nu = 2$ left and right interpolation points, with $\mathbf{r}_1 = \mathbf{r}_2 = 1$ and $\boldsymbol{\ell}_1 = \boldsymbol{\ell}_2 = 1$. We will study various choices of interpolation points, all of which satisfy, for each $\widehat{z} \in \{\lambda_1, \lambda_2, \mu_1, \mu_2\}$,

$$\text{rank}(\widehat{z}\, \mathbb{L} - \mathbb{L}_s) = \text{rank}(\mathbb{L}) = \text{rank}(\mathbb{L}_s) = n = 2. \quad (15)$$

This basic set-up makes it easy to focus on the influence of the interpolation points $\lambda_1, \lambda_2, \mu_1, \mu_2$. We will use the pseudospectra $\sigma_\varepsilon^{(1,1)}(\mathbb{L}_s, \mathbb{L})$ to examine how the interpolation points affect the stability of the eigenvalues of the Loewner pencil.

Table 2 records four different choices of $\{\lambda_1, \lambda_2, \mu_1, \mu_2\}$; Fig. 1 shows the corresponding pseudospectra $\sigma_\varepsilon^{(1,1)}(\mathbb{L}_s, \mathbb{L})$. All four Loewner realizations match the eigenvalues of $\mathbf{A}$ and satisfy the interpolation conditions. However, the pseudospectra show how the *stability* of the eigenvalues -0.1 and -2.1 differs across these

Table 2 Right and left interpolation points for the two-dimensional SISO system (14). The two right columns report the singular values of the Loewner matrix $\mathbb{L}$

| Example	λ_1	λ_2	μ_1	μ_2	$s_1(\mathbb{L})$	$s_2(\mathbb{L})$
(a)	0	1	1i	-1i	6.9871212	0.0731542
(b)	0.25	0.75	2i	-2i	1.0021659	0.0296996
(c)	0.40	0.60	4i	-4i	0.3605151	0.0057490
(d)	8	9	10	11	0.0035344	0.0000019
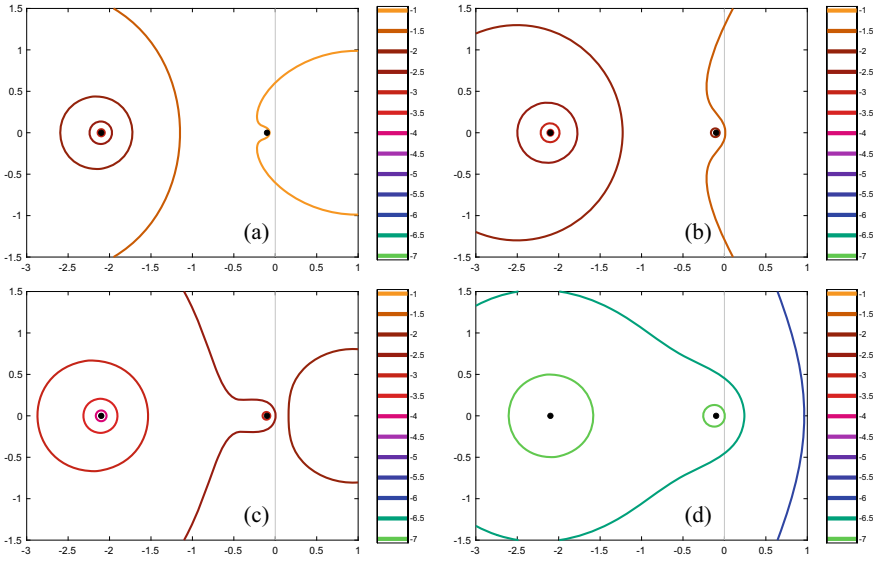
Fig. 1 Boundaries of pseudospectra $\sigma_\varepsilon^{(1,1)}(\mathbb{L}_s, \mathbb{L})$ for four Loewner realizations of the system (14) using the interpolation points in Table 2. All four realizations correctly give $\sigma(\mathbb{L}_s, \mathbb{L}) = \sigma(\mathbf{A})$, but $\sigma_\varepsilon^{(1,1)}(\mathbb{L}_s, \mathbb{L})$ show how the stability of the realized eigenvalues depends on the choice of interpolation points. *In this and all similar plots, the colors denote* $\log_{10}(\varepsilon)$. Thus, in plot (d), there exist perturbations to $\mathbb{L}_s$ and $\mathbb{L}$ of norm $10^{-6.5}$ that move an eigenvalue into the right half-plane

four realizations. From example (a) to (d), these eigenvalues become increasingly sensitive as the interpolation points move farther from $\sigma(\mathbf{A})$. Table 2 also shows the singular values of the Loewner matrix $\mathbb{L}$, demonstrating how the second singular value $s_2(\mathbb{L})$ decreases as the eigenvalues become increasingly sensitive. (Taken to a greater extreme, it would eventually be difficult to determine if $\mathbb{L}$ truly is rank 2.)

Remark 5 By Remark 1, note that if $\varepsilon < s_{\min}(\mathbb{L})$, then $\sigma_\varepsilon^{(1,1)}(\mathbb{L}_s, \mathbb{L})$ will be unbounded. Thus the decreasing values of $s_{\min}(\mathbb{L})$ in Table 2 suggest the enlarging pseudospectra seen in Fig. 1. For example, in case (d) the $\varepsilon = 10^{-5}$ pseudospectrum $\sigma_\varepsilon^{(1,1)}(\mathbb{L}_s, \mathbb{L})$ must contain the point at infinity.

Contrast these results with the standard pseudospectra of $\mathbf{A}$ itself, $\sigma_\varepsilon(\mathbf{A}) = \sigma_\varepsilon^{(1,0)}(\mathbf{A}, \mathbf{I})$ shown at the top of Fig. 2. Since $\mathbf{A}$ is real symmetric (hence normal), $\sigma_\varepsilon(\mathbf{A})$ is the union of open ε-balls surrounding the eigenvalues. Figure 2 compares these pseudospectra to $\sigma_\varepsilon(\mathbb{L}^{-1}\mathbb{L}_s)$, which give insight into the transient behavior of solutions to $\mathbb{L}\dot{\mathbf{x}}(t) = \mathbb{L}_s\mathbf{x}(t)$, e.g., via the bound (11). (Since $\mathbb{L}_s - \mathbb{L}\mathbf{\Lambda} = \mathbf{V}^*\mathbf{R}$, and $m = 1$, $\mathbb{L}^{-1}\mathbb{L}_s = \mathbf{\Lambda} + \mathbb{L}^{-1}\mathbf{V}^*\mathbf{R}$ is a rank-1 update of the matrix $\mathbf{\Lambda}$ [19, p. 643].) The top plot shows $\sigma_\varepsilon(\mathbf{A})$, whose rightmost extent in the complex plane is always $\varepsilon - 0.1$: no transient growth is possible for this system. However, in all four Loewner realizations, $\sigma_\varepsilon(\mathbb{L}^{-1}\mathbb{L}_s)$ extends more than ε into the right-half plane for $\varepsilon = 10^0$
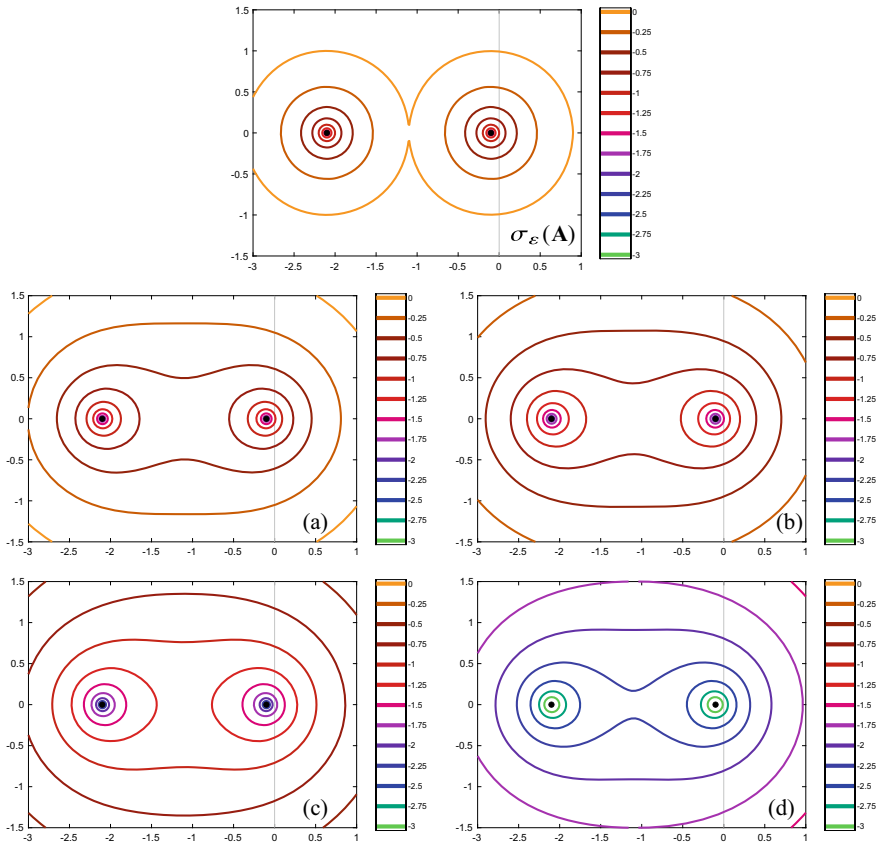
Fig. 2 The pseudospectra $\sigma_\varepsilon(\mathbf{A})$ (top), compared to $\sigma_\varepsilon(\mathbb{L}^{-1}\mathbb{L}_s)$ for four Loewner realizations of the system (14) using the interpolation points in Table 2. In all cases $\sigma(\mathbb{L}^{-1}\mathbb{L}_s) = \sigma(\mathbf{A})$, but the pseudospectra of $\mathbb{L}^{-1}\mathbb{L}_s$ are all quite a bit larger than $\sigma_\varepsilon(\mathbf{A})$

(orange level curve), indicating by (11) that transient growth must occur for some initial condition. Figure 3 shows this growth for all four realizations: *the more remote interpolation points lead to Loewner realizations with greater transient growth.*

6   Example 2: Partitioning Interpolation Points and Noisy Data

To further investigate how the interpolation points influence eigenvalue stability for the Loewner pencil, consider the SISO system of order 10 given by

$$\mathbf{E} = \mathbf{I}, \quad \mathbf{A} = \mathrm{diag}(-1, -2, \ldots, -10), \quad \mathbf{B} = [1, 1, \ldots, 1]^T, \quad \mathbf{C} = [1, 1, \ldots, 1].$$
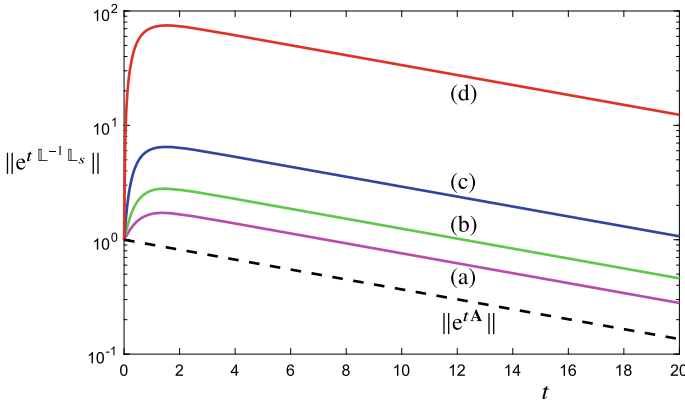
Fig. 3 Evolution of the norm of the solution operator for the original system (black dashed line) and the four interpolating Loewner models. The instability revealed by the pseudospectra in Fig. 2 corresponds to transient growth in the Loewner systems

Figure 4 shows $\sigma_\varepsilon^{(1,1)}(\mathbb{L}_s, \mathbb{L})$ for six configurations of the interpolation points. Plots (a) and (b) use the points $\{-10.25, -9.75, -9.25, \ldots, -1.25, -0.75\}$; in plot (a) the left and right points interleave, suggesting slower decay of the singular values of $\mathbb{L}$, as discussed in Sect. 2.1; in plot (b) the left and right points are separated, leading to faster decay of the singular values of $\mathbb{L}$ and considerably larger pseudospectra. (The Beckermann–Townsend bound [3, Corollary 4.2] applies to this case.) Plot (c) further separates the left and right points, giving even larger pseudospectra. In plots (d) and (f), complex interpolation points $\{-5 \pm 0.5\mathrm{i}, -5 \pm 1.0\mathrm{i}, \ldots, -5 \pm 5\mathrm{i}\}$ are mostly interleaved (d) (keeping conjugate pairs together) and separated (f): the latter significantly enlarges the pseudospectra. Plot (f) uses the same relative arrangement that gave such nice results in plot (a) (the singular value bound (3) is the same for (a) and (e)), but their locations relative to the poles of the original system differ. The pseudospectra are now much larger, showing that a large upper bound in (3) is not alone enough to guarantee small pseudospectra. (Indeed, the pseudospectra are so large in (e) and (f) that the plots are dominated by numerical artifacts of computing $\|(z\mathbb{L} - \mathbb{L}_s)^{-1}\|$.) *Pseudospectra reveal the great influence interpolation point location and partition can have on the stability of the realized pencils.*

Pseudospectra also give insight into the consequences of inexact measurement data. (For a recent study of how noisy data affects the recovered transfer function, see [6].) Consider the following experiment. Take the scenario in Fig. 4a, the most robust of these examples. Subject each right and left measurement $\{\mathbf{w}_1, \ldots, \mathbf{w}_\varrho\} \subset \mathbb{C}$ and $\{\mathbf{v}_1, \ldots, \mathbf{v}_\nu\} \subset \mathbb{C}$ to random complex noise of magnitude 10^{-1}, then build the Loewner pencil $z\widehat{\mathbb{L}} - \widehat{\mathbb{L}}_s$ from this noisy data. How do the badly polluted measurements affect the computed eigenvalues? Figure 5 shows the results of 1,000 random trials, which can depart from the true matrices significantly:

$$3.06 \leq \|\widehat{\mathbb{L}}_s - \mathbb{L}_s\| \leq 5.88, \qquad 0.49 \leq \|\widehat{\mathbb{L}} - \mathbb{L}\| \leq 0.63.$$
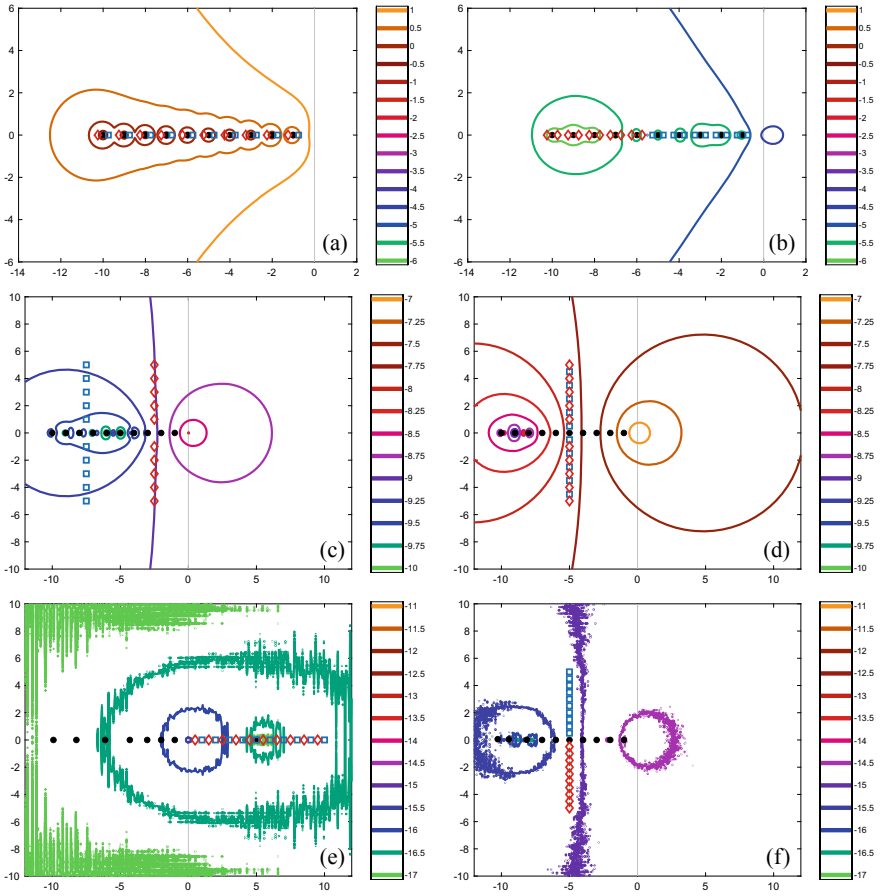
Fig. 4 Pseudospectra $\sigma_\varepsilon^{(1,1)}(\mathbb{L}_s, \mathbb{L})$ for six Loewner realizations of a SISO system of order $n = 10$ with poles $\sigma(\mathbf{A}) = \{-1, -2, \ldots, -10\}$. Black dots show computed eigenvalues of the pencil $z\mathbb{L} - \mathbb{L}_s$ (which should agree with $\sigma(\mathbf{A})$, but a few are off axis in plot (e)); blue squares show the right interpolation points $\{\lambda_i\}$; red diamonds show the left points $\{\mu_j\}$

Despite these large perturbations, the recovered eigenvalues are remarkably accurate: in 99.99% of cases, the eigenvalues have absolute accuracy of at least 10^{-2}, indeed more accurate than the measurements themselves. The pseudospectra in Fig. 4a suggest good robustness (though the pseudospectral level curves are pessimistic by one or two orders of magnitude). Contrast this with the complex interleaved interpolation points used in Fig. 4d. Now we only perturb the data by a small amount, $5 \cdot 10^{-9}$, for which the perturbed Loewner matrices (over 1,000 trials) satisfy

$$1.52 \cdot 10^{-7} \leq \|\widehat{\mathbb{L}}_s - \mathbb{L}_s\| \leq 2.03 \cdot 10^{-7}, \qquad 2.57 \cdot 10^{-8} \leq \|\widehat{\mathbb{L}} - \mathbb{L}\| \leq 3.14 \cdot 10^{-8}.$$
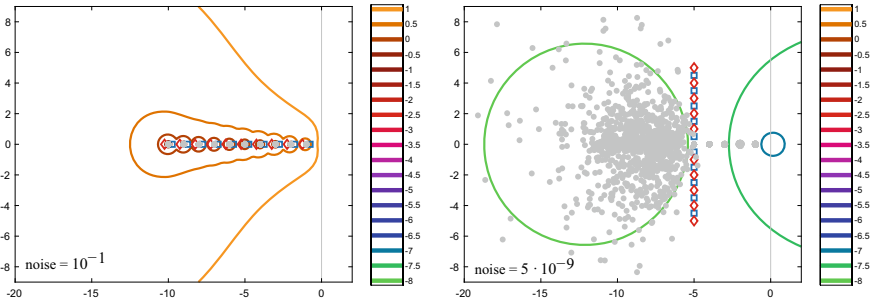
Fig. 5 Eigenvalues of the perturbed Loewner pencil $\widehat{\mathbb{L}}_s - z\widehat{\mathbb{L}}$ (gray dots), constructed from measurements that have been perturbed by random complex noise of magnitude 10^{-1} (left) and $5 \cdot 10^{-9}$ (right) (1,000 trials). As the pseudospectra in Fig. 4a, d indicate, the interleaved interpolation points on the left are remarkably stable, while the similarly interleaved complex interpolation points on the right give a Loewner pencil that is highly sensitive to small changes in the data

With this mere hint of noise, the eigenvalues of the recovered system erupt: only 36.73% of the eigenvalues are correct to two digits. (Curiously, -4 and -5 are always computed correctly, while -8, -9, and -10 are never computed correctly.) The pseudospectra indicate that the leftmost eigenvalues are more sensitive, and again hint at the effect of the perturbation (though off by roughly an order of magnitude in ε).[3] Measurements of real systems (or even numerical simulations of nontrivial systems) are unlikely to produce such high accuracy; pseudospectra can reveal the virtue or folly of a given interpolation point configuration.

In these simple experiments, pseudospectra have been most helpful for indicating the sensitivity of eigenvalues when the left and right interpolation points are favorably
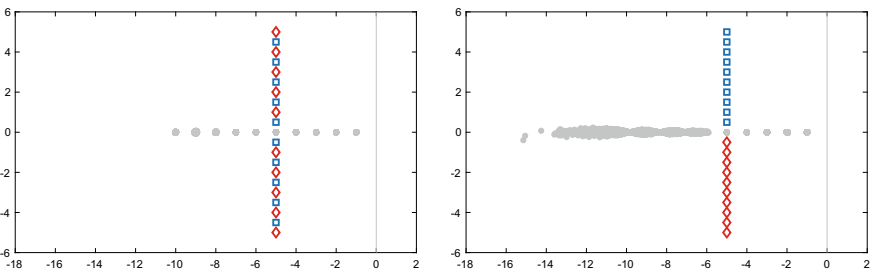
Fig. 6 Eigenvalues of the perturbed Loewner pencil $\widehat{\mathbb{L}}_s - z\widehat{\mathbb{L}}$ (gray dots), constructed from measurements that have been perturbed by random complex noise of magnitude 10^{-10} (10,000 trials). As suggested by the pseudospectra plots, the interleaved interpolation points (left) are more robust to perturbations than the separated points (right), though the difference is not as acute as suggested by Fig. 4d, f. For example, in these 10,000 trials, the least stable pole (-9) is computed accurately (absolute error less than 0.01) in 10.97% of trials on the left, and 0.29% on the right

---

[3] One might also consider the $\varepsilon = 1$ level curve of $\sigma_\varepsilon^{(\gamma_\star, \delta_\star)}(\mathbb{L}_s, \mathbb{L})$ for $\gamma_\star = \max\{\|\mathbb{L}_s - \widehat{\mathbb{L}}_s\|\}$ and $\delta_\star = \max\{\|\mathbb{L} - \widehat{\mathbb{L}}\|\}$.

partitioned (e.g., interleaved). They seem to be less precise at predicting the sensitivity to noise of poor left/right partitions of the interpolation points. Figure 6 gives an example, based on the two partitions of the same interpolation points in Fig. 4d, f. The pseudospectra suggest that the eigenvalues for plot (f) should be much more sensitive to noise than those for the interleaved points in plot (d). In fact, the configuration in plot (f) appears to be only marginally less stable to noise of size 10^{-10}, over 10,000 trials. This is a case where one could potentially glean additional insight from *structured* Loewner pseudospectra.

7   Conclusion

Pseudospectra provide a tool for analyzing the stability of eigenvalues of Loewner matrix pencils. Elementary examples show how pseudospectra can inform the selection and partition of interpolation points, and bound the eigenvalues of Loewner pencils in the presence of noisy data. Using a different approach to pseudospectra, we showed that while the realized Loewner pencil matches the poles of the original system, it need not replicate transient dynamics of $\mathbf{E}\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t)$; pseudospectra can reveal potential transient growth, which varies with the interpolation points.

In this initial study we have intentionally used simple examples involving small, symmetric $\mathbf{A}$ and $\mathbf{E} = \mathbf{I}$. Realistic examples, e.g., with complex poles, nonnormal $\mathbf{A}$, singular $\mathbf{E}$, multiple inputs and outputs, and rank-deficient Loewner matrices, will add additional complexity. Moreover, we have only sought to realize a system whose order is known; we have not addressed pseudospectra of the reduced pencils (7) in the context of data-driven model reduction.

Structured Loewner pseudospectra provide another avenue for future study. Structured *matrix pencil* pseudospectra have not been much investigated, especially with Loewner structure. Rump's results for standard pseudospectra [25] suggest the following problem; its positive resolution would imply that the Loewner pseudospectrum $\sigma_\varepsilon^{(\gamma,\delta)}(\mathbb{L}_s, \mathbb{L})$ matches the structured Loewner matrix pseudospectrum.

Given any $\varepsilon, \gamma, \delta > 0$, Loewner matrix $\mathbb{L}$ and associated shifted Loewner matrix $\mathbb{L}_s$, suppose $z \in \sigma_\varepsilon^{(\gamma,\delta)}(\mathbb{L}_s, \mathbb{L})$. *Does there exist some Loewner matrix $\widehat{\mathbb{L}}$ and associated shifted Loewner matrix $\widehat{\mathbb{L}}_s$ such that $z \in \sigma(\widehat{\mathbb{L}}_s, \widehat{\mathbb{L}})$ and*

$$\|\widehat{\mathbb{L}}_s - \mathbb{L}_s\| < \epsilon\gamma, \qquad \|\widehat{\mathbb{L}} - \mathbb{L}\| < \epsilon\delta \ \ ?$$

Appendix

We provide a MATLAB implementation that computes the inverse iteration vectors **u** in (12) in only $\mathcal{O}(n^2)$ operations by exploiting the Cauchy-like rank displacement structure of the Loewner pencil, as shown in (13). Namely, we start from the general $\mathcal{O}(n^3)$ MATLAB code from [29, p. 373] and modify it to account for the Loewner structure, and hence achieve $\mathcal{O}(n^2)$ efficiency. This code computes $\|(z\mathbb{L} - \mathbb{L}_s)^{-1}\|$ for a fixed z. To compute pseudospectra, one applies this algorithm on a grid of z values. In that case, the $\mathcal{O}((m + p)n^2)$ structured LU factorization in the first line need only be computed *once* for all z values (just as the standard algorithm computes an $\mathcal{O}(n^3)$ simultaneous triangularization using the QZ algorithm once for all z).

```
[L1,U1,piv1] = LUdispPiv(mu,lambda,[V' L'],[R.' -W.']);
L1t = L1'; U1t = U1';
z = 1./(z-lambda); Upz = z.*(U1\(L1\V(:,piv1)'));
[L2,U2,piv2] = lu(eye(m)-R*Upz,'vector'); R2 = R(piv2,:);
L2t = L2'; U2t = U2'; R2t = R2'; Upzt = Upz';
applyTheta = @(x) x+Upz*(U2\(L2\(R2*x)));
applyThetaTranspose = @(x) x+R2t*(L2t\(U2t\(Upzt*x)));

sigold = 0;
for it = 1:maxit
    u = U1\(L1\u(piv1));
    u = conj(z).*applyThetaTranspose(applyTheta(z.*u));
    u(piv1) = L1t\(U1t\u);
    sig = 1/norm(u);
    if abs(sigold/sig-1) < 1e-2, break, end
    u = sig*u; sigold = sig;
end
sigmin = sqrt(sig);
```

The function `LUdispPiv` computes the LU factorization (with partial pivoting) of the Loewner matrix $\mathbb{L}$ in $\mathcal{O}((m + p)n^2)$ operations. The Loewner matrix is not formed explicitly; instead, the function uses the raw interpolation data λ_i, $\mathbf{r}_i$, $\mathbf{w}_i$ and μ_j, $\boldsymbol{\ell}_j$, $\mathbf{v}_j$. The implementation details for `LUdispPiv` can be found in [12, Sect. 12.1]. The LU factorization of $\mathbf{I} - \mathbf{RV}^*(z) \in \mathbb{C}^{m \times m}$ is given by `L2` and `U2`, while the first three lines of the loop represent the computation of $\mathbb{L}^{-*}\boldsymbol{\Theta}(z)^*\boldsymbol{\Theta}(z)\mathbb{L}^{-1}\mathbf{u}$, as defined in (13). Note the careful grouping of terms and the use of elementwise multiplication `.*` to keep the total operation count at $\mathcal{O}(n^2)$.

References

1. Antoulas, A.C., Lefteriu, S., Ionita, A.C.: A tutorial introduction to the Loewner framework for model reduction. In: Benner, P., Cohen, A., Ohlberger, M., Willcox, K. (eds.) Model Reduction and Approximation: Theory and Algorithms, pp. 335–376. SIAM, Philadelphia (2017)
2. Antoulas, A.C., Sorensen, D.C., Zhou, Y.: On the decay rate of Hankel singular values and related issues. Syst. Control Lett. **46**, 323–342 (2002)
3. Beckermann, B., Townsend, A.: On the singular values of matrices with displacement structure. SIAM J. Matrix Anal. Appl. **38**, 1227–1248 (2017)
4. Chaitin-Chatelin, F., Fraysse, V.: Lectures on Finite Precision Computations. SIAM, Philadelphia (1996)
5. Chaitin-Chatelin, F., Harrabi, A.: About definitions of pseudospectra of closed operators in Banach spaces. Technical Report TR/PA/98/08, CERFACS (1998)
6. Drmač, Z., Peherstorfer, B.: Learning low-dimensional dynamical-system models from noisy frequency-response data with Loewner rational interpolation. In: C. Beattie et al. (eds.) Realization and Model Reduction of Dynamical Systems, pp. 39–57. Springer, Cham (2022)
7. Embree, M.: Pseudospectra (Chap. 23). In: Hogben, L. (ed.) Handbook of Linear Algebra, 2nd edn. CRC/Taylor & Francis, Boca Raton (2014)
8. Embree, M., Keeler, B.: Pseudospectra of matrix pencils for transient analysis of differential-algebraic equations. SIAM J. Matrix Anal. Appl. **38**, 1028–1054 (2017)
9. Filip, S.-I., Nakatsukasa, Y., Trefethen, L.N., Beckermann, B.: Rational minimax approximation via adaptive barycentric approximations. SIAM J. Sci. Comput. **40**, A2427–A2455 (2018)
10. Fraysse, V., Gueury, M., Nicoud, F., Toumazou, V.: Spectral portraits for matrix pencils. Technical Report TR/PA/96/19, CERFACS (1996)
11. Gohberg, I., Kailath, T., Olshevsky, V.: Fast Gaussian elimination with partial pivoting for matrices with displacement structure. Math. Comp. **64**, 1557–1576 (1995)
12. Golub, G.H., Van Loan, C.F.: Matrix Computations, 4th edn. Johns Hopkins University Press, Baltimore (2012)
13. Gosea, I.V., Antoulas, A.C.: Rational approximation of the absolute value function from measurements: a numerical study of recent methods (2020). arXiv:2005.02736
14. Higham, N.J., Tisseur, F.: More on pseudospectra for polynomial eigenvalue problems and applications in control theory. Linear Algebra Appl. **351–352**, 435–453 (2002)
15. Hinrichsen, D., Pritchard, A.J.: Real and complex stability radii: a survey. In: Hinrichsen, D., Mårtensson, B. (eds.) Control of Uncertain Systems. Birkhäuser, Boston (1990)
16. Horn, R.A., Johnson, C.R.: Matrix Analysis, 2nd edn. Cambridge University Press, Cambridge (2013)
17. Ioniă, A.C.: Lagrange rational interpolation and its applications to approximation of large-scale dynamical systems. Ph.D. Thesis, Rice University (2013)
18. Lavallée, P.-F.: Nouvelles Approches de Calcul du ε-Spectre de Matrices et de Faisceaux de Matrices. Ph.D. Thesis, Université de Rennes 1 (1997)
19. Mayo, A.J., Antoulas, A.C.: A framework for the solution of the generalized realization problem. Linear Algebra Appl. **425**, 634–662 (2007)
20. Overton, M.L.: Numerical Computing with IEEE Floating Point Arithmetic. SIAM, Philadelphia (2001)
21. Penzl, T.: A cyclic low-rank Smith method for large sparse Lyapunov equations. SIAM J. Sci. Comput. **21**, 1401–1418 (2000)
22. Penzl, T.: Eigenvalue decay bounds for solutions of Lyapunov equations: the symmetric case. Syst. Control Lett. **40**, 139–144 (2000)
23. Riedel, K.S.: Generalized epsilon-pseudospectra. SIAM J. Numer. Anal. **31**, 1219–1225 (1994)
24. Ruhe, A.: The rational Krylov algorithm for large nonsymmetric eigenvalues — mapping the resolvent norms (pseudospectrum). Unpublished manuscript, March 1995
25. Rump, S.M.: Eigenvalues, pseudospectrum and structured perturbations. Linear Algebra Appl. **413**, 567–593 (2006)

26. Sabino, J.: Solution of large-scale Lyapunov equations via the block modified Smith method. Ph.D. Thesis, Rice University (2006)
27. Tisseur, F., Higham, N.J.: Structured pseudospectra for polynomial eigenvalue problems, with applications. SIAM J. Matrix Anal. Appl. **23**, 187–208 (2001)
28. Trefethen, L.N.: Computation of pseudospectra. Acta Numer. **8**, 247–295 (1999)
29. Trefethen, L.N., Embree, M.: Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators. Princeton University Press, Princeton (2005)
30. Trefethen, L.N., Trefethen, A.E., Reddy, S.C., Driscoll, T.A.: Hydrodynamic stability without eigenvalues. Science **261**, 578–584 (1993)
31. Wright, T.G.: Algorithms and software for pseudospectra. D.Phil. Thesis, Oxford University (2002)