




Pseudospectra and Nonnormal Dynamical Systems

Mark Embree and Russell Carden
Computational and Applied Mathematics
Rice University
Houston, Texas

ELGERSBURG

MARCH 2012

Overview of the Course

*These lectures describe modern tools
for the spectral analysis of dynamical systems.*

We shall cover a mix of theory, computation, and applications.

By the end of the week, you will have a thorough understanding of the sources of nonnormality, its affects on the behavior of dynamical systems, tools for assessing its impact, and examples where it arises.

You will be able to understand phenomena that many people find quite mysterious when encountered in the wild.

Overview of the Course

*These lectures describe modern tools
for the spectral analysis of dynamical systems.*

We shall cover a mix of theory, computation, and applications.

By the end of the week, you will have a thorough understanding of the sources of nonnormality, its effects on the behavior of dynamical systems, tools for assessing its impact, and examples where it arises.

You will be able to understand phenomena that many people find quite mysterious when encountered in the wild.

Lecture 1: Introduction to Nonnormality and Pseudospectra

Lecture 2: Functions of Matrices

Lecture 3: Toeplitz Matrices and Model Reduction

Lecture 4: Model Reduction, Numerical Algorithms, Differential Operators

Lecture 5: Discretization, Extensions, Applications

Lecture 1: Introduction to Nonnormality and Pseudospectra

- ▶ Some motivating examples
- ▶ Normality and nonnormality
- ▶ Numerical range (field of values)
- ▶ Pseudospectra
- ▶ Computing the numerical range and pseudospectra, EigTool

Notation

$\mathbf{A}, \mathbf{B}, \dots \in \mathbf{C}^{m \times n}$	$m \times n$ matrices with complex entries
$\mathbf{x}, \mathbf{y}, \dots \in \mathbf{C}^n$	column vectors with n complex entries
$\mathbf{A}^*$	conjugate transpose, $\mathbf{A}^* = \overline{\mathbf{A}}^T \in \mathbf{C}^{n \times m}$
$\mathbf{x}^*$	conjugate transpose (row vector), $\mathbf{x}^* = \overline{\mathbf{x}}^T$
$\dot{\mathbf{x}}(t)$	time-derivative of a vector-valued function $\mathbf{x} : \mathbf{R} \rightarrow \mathbf{C}^n$.
$\ \mathbf{x}\ $	norm of $\mathbf{x}$ (generally the Euclidean norm, $\ \mathbf{x}\ = \sqrt{\mathbf{x}^* \mathbf{x}}$)
$\ \mathbf{A}\ $	matrix norm of $\mathbf{A}$, induced by the vector norm: $\ \mathbf{A}\ := \max_{\ \mathbf{x}\ =1} \ \mathbf{A}\mathbf{x}\ $ $\ \mathbf{A}\mathbf{B}\ \leq \ \mathbf{A}\ \ \mathbf{B}\ $ (submultiplicativity)

Notation

$\mathbf{A}, \mathbf{B}, \dots \in \mathbf{C}^{m \times n}$	$m \times n$ matrices with complex entries
$\mathbf{x}, \mathbf{y}, \dots \in \mathbf{C}^n$	column vectors with n complex entries
$\mathbf{A}^*$	conjugate transpose, $\mathbf{A}^* = \overline{\mathbf{A}}^T \in \mathbf{C}^{n \times m}$
$\mathbf{x}^*$	conjugate transpose (row vector), $\mathbf{x}^* = \overline{\mathbf{x}}^T$
$\dot{\mathbf{x}}(t)$	time-derivative of a vector-valued function $\mathbf{x} : \mathbf{R} \rightarrow \mathbf{C}^n$.
$\ \mathbf{x}\ $	norm of $\mathbf{x}$ (generally the Euclidean norm, $\ \mathbf{x}\ = \sqrt{\mathbf{x}^* \mathbf{x}}$)
$\ \mathbf{A}\ $	matrix norm of $\mathbf{A}$, induced by the vector norm: $\ \mathbf{A}\ := \max_{\ \mathbf{x}\ =1} \ \mathbf{A}\mathbf{x}\ $ $\ \mathbf{A}\mathbf{B}\ \leq \ \mathbf{A}\ \ \mathbf{B}\ $ (submultiplicativity)
$\sigma(\mathbf{A})$	spectrum (eigenvalues) of $\mathbf{A} \in \mathbf{C}^{n \times n}$: $\sigma(\mathbf{A}) = \{z \in \mathbf{C} : z\mathbf{I} - \mathbf{A} \text{ is not invertible}\}$
$s_k(\mathbf{A})$	k th largest singular value of $\mathbf{A} \in \mathbf{C}^{m \times n}$
$s_{\min}(\mathbf{A})$	smallest singular value of $\mathbf{A} \in \mathbf{C}^{m \times n}$
$\text{Ran}(\mathbf{V})$	range (column space) of the matrix $\mathbf{V} \in \mathbf{C}^{n \times k}$
$\text{Ker}(\mathbf{V})$	kernel (null space) of the matrix $\mathbf{V} \in \mathbf{C}^{n \times k}$

Basic Stability Theory

We begin with a standard time-invariant linear system

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t)$$

with given initial state $\mathbf{x}(0) = \mathbf{x}_0$.

This system has the well-known solution

$$\mathbf{x}(t) = e^{t\mathbf{A}}\mathbf{x}_0,$$

where $e^{t\mathbf{A}}$ is the *matrix exponential*.

Basic Stability Theory

We begin with a standard time-invariant linear system

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t)$$

with given initial state $\mathbf{x}(0) = \mathbf{x}_0$.

This system has the well-known solution

$$\mathbf{x}(t) = e^{t\mathbf{A}}\mathbf{x}_0,$$

where $e^{t\mathbf{A}}$ is the *matrix exponential*.

For $\mathbf{A}$ is diagonalizable,

$$\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1} = \begin{bmatrix} \mathbf{v}_1 & \mathbf{v}_2 & \cdots & \mathbf{v}_n \end{bmatrix} \begin{bmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \ddots & \\ & & & \lambda_n \end{bmatrix} \begin{bmatrix} \hat{\mathbf{v}}_1^* \\ \hat{\mathbf{v}}_2^* \\ \vdots \\ \hat{\mathbf{v}}_n^* \end{bmatrix} = \sum_{j=1}^n \lambda_j \mathbf{v}_j \hat{\mathbf{v}}_j^*,$$
$$e^{t\mathbf{A}} = \begin{bmatrix} \mathbf{v}_1 & \mathbf{v}_2 & \cdots & \mathbf{v}_n \end{bmatrix} \begin{bmatrix} e^{t\lambda_1} & & & \\ & e^{t\lambda_2} & & \\ & & \ddots & \\ & & & e^{t\lambda_n} \end{bmatrix} \begin{bmatrix} \hat{\mathbf{v}}_1^* \\ \hat{\mathbf{v}}_2^* \\ \vdots \\ \hat{\mathbf{v}}_n^* \end{bmatrix} = \sum_{j=1}^n e^{t\lambda_j} \mathbf{v}_j \hat{\mathbf{v}}_j^*.$$

Basic Stability Theory

The system $\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t)$ is said to be *stable* provided the solution

$$\mathbf{x}(t) = e^{t\mathbf{A}}\mathbf{x}_0$$

decays as $t \rightarrow \infty$, i.e., $\|\mathbf{x}(t)\| \rightarrow 0$ for all initial states $\mathbf{x}_0$.

Recall that (for diagonalizable $\mathbf{A}$)

$$e^{t\mathbf{A}} = \begin{bmatrix} \mathbf{v}_1 & \mathbf{v}_2 & \cdots & \mathbf{v}_n \end{bmatrix} \begin{bmatrix} e^{t\lambda_1} & & & \\ & e^{t\lambda_2} & & \\ & & \ddots & \\ & & & e^{t\lambda_n} \end{bmatrix} \begin{bmatrix} \hat{\mathbf{v}}_1^* \\ \hat{\mathbf{v}}_2^* \\ \vdots \\ \hat{\mathbf{v}}_n^* \end{bmatrix} = \sum_{j=1}^n e^{t\lambda_j} \mathbf{v}_j \hat{\mathbf{v}}_j^*.$$

Similar formulas hold for nondiagonalizable $\mathbf{A}$.

Since, e.g.,

$$\begin{aligned} \|\mathbf{x}(t)\| &\leq \|e^{t\mathbf{A}}\| \|\mathbf{x}_0\| \\ &\leq \|\mathbf{V}\| \|\mathbf{V}^{-1}\| \max_{\lambda \in \sigma(\mathbf{A})} |e^{t\lambda}| \leq \|\mathbf{V}\| \|\mathbf{V}^{-1}\| \max_{\lambda \in \sigma(\mathbf{A})} e^{t \operatorname{Re} \lambda}, \end{aligned}$$

we say that the system (or $\mathbf{A}$) is *stable* provided *all eigenvalues of $\mathbf{A}$ are in the left half of the complex plane.*


1(a) Some Motivating Examples

Quiz

The plots below show the solution to two dynamical systems, $\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t)$.

Which system is stable?

That is, for which system, does $\mathbf{x}(t) \rightarrow \mathbf{0}$ as $t \rightarrow \infty$?




- (a) neither system is stable
- (b) only the one on the blue system is stable
- (c) only the one on the red system is stable
- (d) both are stable

Quiz

The plots below show the solution to two dynamical systems, $\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t)$.

Which system is stable?

That is, for which system, does $\mathbf{x}(t) \rightarrow \mathbf{0}$ as $t \rightarrow \infty$?




- (a) neither system is stable
- (b) only the one on the blue system is stable
- (c) only the one on the red system is stable
- (d) both are stable

Quiz


The plots below show the solution to two dynamical systems, $\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t)$.

Which system is stable?

That is, for which system, does $\mathbf{x}(t) \rightarrow \mathbf{0}$ as $t \rightarrow \infty$?



original unstable system



stabilized system

Eigenvalues of 55×55 Boeing 767 flutter models [Burke, Lewis, Overton].

These eigenvalues do not reveal the exotic transient behavior.

Why is Transient Growth Important?

Many linear systems arise from the linearization of nonlinear equations, e.g., Navier–Stokes. We compute eigenvalues as part of *linear stability analysis*.

Transient growth in a stable linearized system has implications for the behavior of the associated nonlinear system.

Seminal article in *Science*, 1993:

Hydrodynamic Stability Without Eigenvalues

Lloyd N. Trefethen, Anne E. Trefethen, Satish C. Reddy, Tobin A. Driscoll

Fluid flows that are smooth at low speeds become unstable and then turbulent at higher speeds. This phenomenon has traditionally been investigated by linearizing the equations of flow and testing for unstable eigenvalues of the linearized problem, but the results of such investigations agree poorly in many cases with experiments. Nevertheless, linear effects play a central role in hydrodynamic instability. A reconciliation of these findings with the traditional analysis is presented based on the “pseudospectra” of the linearized problem, which imply that small perturbations to the smooth flow may be amplified by factors on the order of 10^5 by a linear mechanism even though all the eigenmodes decay monotonically. The methods suggested here apply also to other problems in the mathematical sciences that involve nonorthogonal eigenfunctions.

Why is Transient Growth Important?

Transient growth in a stable linearized system has implications for the behavior of the associated nonlinear system.

(RECIPE FOR LINEAR STABILITY ANALYSIS)

Consider the autonomous nonlinear system $\dot{\mathbf{u}}(t) = \mathbf{f}(\mathbf{u})$.

- ▶ Find a steady state $\mathbf{u}_*$, i.e., $\mathbf{f}(\mathbf{u}_*) = \mathbf{0}$.
- ▶ Linearize $\mathbf{f}$ about this steady state and analyze small perturbations, $\mathbf{u} = \mathbf{u}_* + \mathbf{v}$:

$$\begin{aligned}\dot{\mathbf{v}}(t) = \dot{\mathbf{u}}(t) &= \mathbf{f}(\mathbf{u}_* + \mathbf{v}) \\ &= \mathbf{f}(\mathbf{u}_*) + \mathbf{A}\mathbf{v} + \mathcal{O}(\|\mathbf{v}\|^2) \\ &= \mathbf{A}\mathbf{v} + \mathcal{O}(\|\mathbf{v}\|^2).\end{aligned}$$

- ▶ Ignore higher-order effects, and analyze the linear system $\dot{\mathbf{v}}(t) = \mathbf{A}\mathbf{v}(t)$. The steady state $\mathbf{u}_*$ is *stable* provided $\mathbf{A}$ is stable.

Why is Transient Growth Important?

Transient growth in a stable linearized system has implications for the behavior of the associated nonlinear system.

(RECIPE FOR LINEAR STABILITY ANALYSIS)

Consider the autonomous nonlinear system $\dot{\mathbf{u}}(t) = \mathbf{f}(\mathbf{u})$.

- ▶ Find a steady state $\mathbf{u}_*$, i.e., $\mathbf{f}(\mathbf{u}_*) = \mathbf{0}$.
- ▶ Linearize $\mathbf{f}$ about this steady state and analyze small perturbations, $\mathbf{u} = \mathbf{u}_* + \mathbf{v}$:

$$\begin{aligned}\dot{\mathbf{v}}(t) = \dot{\mathbf{u}}(t) &= \mathbf{f}(\mathbf{u}_* + \mathbf{v}) \\ &= \mathbf{f}(\mathbf{u}_*) + \mathbf{A}\mathbf{v} + \mathcal{O}(\|\mathbf{v}\|^2) \\ &= \mathbf{A}\mathbf{v} + \mathcal{O}(\|\mathbf{v}\|^2).\end{aligned}$$

- ▶ Ignore higher-order effects, and analyze the linear system $\dot{\mathbf{v}}(t) = \mathbf{A}\mathbf{v}(t)$. The steady state $\mathbf{u}_*$ is *stable* provided $\mathbf{A}$ is stable.

But what if the small perturbation $\mathbf{v}(t)$ grows by orders of magnitude before eventually decaying?

Example: a Nonlinear Heat Equation

An example of the failure of linear stability analysis for a PDE problem.

Consider the nonlinear heat equation on $x \in [-1, 1]$ with $u(-1, t) = u(1, t) = 0$

$$u_t(x, t) = \nu u_{xx}(x, t)$$

with $\nu > 0$

Example: a Nonlinear Heat Equation

An example of the failure of linear stability analysis for a PDE problem.

Consider the nonlinear heat equation on $x \in [-1, 1]$ with $u(-1, t) = u(1, t) = 0$

$$u_t(x, t) = \nu u_{xx}(x, t) + \sqrt{\nu} u_x(x, t)$$

with $\nu > 0$

Example: a Nonlinear Heat Equation

An example of the failure of linear stability analysis for a PDE problem.

Consider the nonlinear heat equation on $x \in [-1, 1]$ with $u(-1, t) = u(1, t) = 0$

$$u_t(x, t) = \nu u_{xx}(x, t) + \sqrt{\nu} u_x(x, t) + \frac{1}{8} u(x, t)$$

with $\nu > 0$

Example: a Nonlinear Heat Equation

An example of the failure of linear stability analysis for a PDE problem.

Consider the nonlinear heat equation on $x \in [-1, 1]$ with $u(-1, t) = u(1, t) = 0$

$$u_t(x, t) = \nu u_{xx}(x, t) + \sqrt{\nu} u_x(x, t) + \frac{1}{8} u(x, t) + u(x, t)^p$$

with $\nu > 0$ and $p > 1$ [Demanet, Holmer, Zworski].

Example: a Nonlinear Heat Equation

An example of the failure of linear stability analysis for a PDE problem.

Consider the nonlinear heat equation on $x \in [-1, 1]$ with $u(-1, t) = u(1, t) = 0$

$$u_t(x, t) = \nu u_{xx}(x, t) + \sqrt{\nu} u_x(x, t) + \frac{1}{8} u(x, t) + u(x, t)^p$$

with $\nu > 0$ and $p > 1$ [Demanet, Holmer, Zworski].

The linearization L , an advection–diffusion operator,

$$Lu = \nu u_{xx} + \sqrt{\nu} u_x + \frac{1}{8} u$$

has eigenvalues and eigenfunctions


$$\lambda_n = -\frac{1}{8} - \frac{n^2 \pi^2 \nu}{4} < 0, \quad u_n(x) = e^{-x/(2\sqrt{\nu})} \sin(n\pi x/2);$$

see, e.g., [Reddy & Trefethen 1994].

The linearized operator is stable for all $\nu > 0$, but has interesting transients

Nonnormality in the Linearization


The linearized system is stable:



Rightmost part of the spectrum for $\nu = 0.002$


But transient growth can feed the nonlinearity...

Evolution of a Small Initial Condition



Nonlinear model (blue) and linearization (black)

Transient Behavior




Linearized system (black) and nonlinear system (dashed blue)

Nonnormal growth feeds the nonlinear instability.

Spectra and Pseudospectra

Source for much of the content of these lectures:



Princeton University Press
2005

Motivating Applications

Situations like these arise in many applications:

- ▶ convective fluid flows
- ▶ damped mechanical systems
- ▶ atmospheric science
- ▶ magnetohydrodynamics
- ▶ neutron transport
- ▶ population dynamics
- ▶ food webs
- ▶ directed social networks
- ▶ Markov chains
- ▶ lasers

$$\mathbf{u}_t(\mathbf{x}, t) = \nu \Delta \mathbf{u}(\mathbf{x}, t)$$

$$\mathbf{M}\ddot{\mathbf{x}}(t) = -\mathbf{K}\mathbf{x}(t)$$

Motivating Applications

Situations like these arise in many applications:

- ▶ convective fluid flows
- ▶ damped mechanical systems
- ▶ atmospheric science
- ▶ magnetohydrodynamics
- ▶ neutron transport
- ▶ population dynamics
- ▶ food webs
- ▶ directed social networks
- ▶ Markov chains
- ▶ lasers

$$\mathbf{u}_t(\mathbf{x}, t) = \nu \Delta \mathbf{u}(\mathbf{x}, t) - (\mathbf{a} \cdot \nabla) \mathbf{u}(\mathbf{x}, t)$$

$$\mathbf{M}\ddot{\mathbf{x}}(t) = -\mathbf{K}\mathbf{x}(t) - \mathbf{D}\dot{\mathbf{x}}(t)$$

1(b) Normality and Nonnormality

Normality and Nonnormality

Unless otherwise noted, all matrices are of size $n \times n$, with complex entries.

The *adjoint* is denoted by $\mathbf{A}^* = \overline{\mathbf{A}^T}$.

Definition (Normal)

The matrix $\mathbf{A}$ is *normal* if it commutes with its adjoint, $\mathbf{A}^* \mathbf{A} = \mathbf{A} \mathbf{A}^*$.

$$\mathbf{A} = \begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix} : \mathbf{A}^* \mathbf{A} = \mathbf{A} \mathbf{A}^* = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix} \implies \text{normal}$$

$$\mathbf{A} = \begin{bmatrix} -1 & 1 \\ 0 & 1 \end{bmatrix} : \mathbf{A}^* \mathbf{A} = \begin{bmatrix} 1 & -1 \\ -1 & 2 \end{bmatrix} \neq \begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix} = \mathbf{A} \mathbf{A}^* \implies \text{nonnormal}$$

Normality and Nonnormality

Unless otherwise noted, all matrices are of size $n \times n$, with complex entries.

The *adjoint* is denoted by $\mathbf{A}^* = \overline{\mathbf{A}^T}$.

Definition (Normal)

The matrix $\mathbf{A}$ is *normal* if it commutes with its adjoint, $\mathbf{A}^* \mathbf{A} = \mathbf{A} \mathbf{A}^*$.

$$\mathbf{A} = \begin{bmatrix} 1 & 1 \\ -1 & 1 \end{bmatrix} : \mathbf{A}^* \mathbf{A} = \mathbf{A} \mathbf{A}^* = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix} \Rightarrow \text{normal}$$

$$\mathbf{A} = \begin{bmatrix} -1 & 1 \\ 0 & 1 \end{bmatrix} : \mathbf{A}^* \mathbf{A} = \begin{bmatrix} 1 & -1 \\ -1 & 2 \end{bmatrix} \neq \begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix} = \mathbf{A} \mathbf{A}^* \Rightarrow \text{nonnormal}$$

Important note:

The adjoint is defined via the inner product: $\langle \mathbf{A}\mathbf{x}, \mathbf{y} \rangle = \langle \mathbf{x}, \mathbf{A}^*\mathbf{y} \rangle$.

hence the definition of normality depends on the inner product.

Here we always use the standard Euclidean inner product, unless noted.

In applications, one *must* use the physically relevant inner product.

Conditions for Normality

Many (~ 89) equivalent definitions of normality are known; see [Grone et al. 1987], [Elsner & Ikramov 1998].

By far, the most important of these concerns the eigenvectors of $\mathbf{A}$.

Theorem

The matrix $\mathbf{A}$ is normal if and only if it is unitarily diagonalizable,

$$\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^*,$$

for $\mathbf{U}$ unitary ($\mathbf{U}^\mathbf{U} = \mathbf{I}$) and $\mathbf{\Lambda}$ diagonal.*

Equivalently, $\mathbf{A}$ is normal if and only if it possesses an orthonormal basis of eigenvectors (i.e., the columns of $\mathbf{U}$).

Conditions for Normality

Many (~ 89) equivalent definitions of normality are known; see [Grone et al. 1987], [Elsner & Ikramov 1998].

By far, the most important of these concerns the eigenvectors of $\mathbf{A}$.

Theorem

The matrix $\mathbf{A}$ is normal if and only if it is unitarily diagonalizable,

$$\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^*,$$

for $\mathbf{U}$ unitary ($\mathbf{U}^\mathbf{U} = \mathbf{I}$) and $\mathbf{\Lambda}$ diagonal.*

Equivalently, $\mathbf{A}$ is normal if and only if it possesses an orthonormal basis of eigenvectors (i.e., the columns of $\mathbf{U}$).

Hence, any nondiagonalizable (defective) matrix is nonnormal.

But there are many interesting diagonalizable nonnormal matrices.

Our fixation with diagonalizability has caused us to overlook these matrices.

Orthogonality of Eigenvectors

Theorem

The matrix $\mathbf{A}$ is normal if and only if it is unitarily diagonalizable,

$$\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^*,$$

for $\mathbf{U}$ unitary ($\mathbf{U}^\mathbf{U} = \mathbf{I}$) and $\mathbf{\Lambda}$ diagonal.*

Equivalently, $\mathbf{A}$ is normal if and only if it possesses an orthonormal basis of eigenvectors (i.e., the columns of $\mathbf{U}$).

An orthogonal basis of eigenvectors gives a perfect coordinate system for studying dynamical systems: set $\mathbf{z}(t) := \mathbf{U}^*\mathbf{x}(t)$, so

$$\begin{aligned}\mathbf{x}'(t) = \mathbf{A}\mathbf{x}(t) &\implies \mathbf{U}^*\mathbf{x}'(t) = \mathbf{U}^*\mathbf{A}\mathbf{U}\mathbf{U}^*\mathbf{x}(t) \\ &\implies \mathbf{z}'(t) = \mathbf{\Lambda}\mathbf{z}(t) \\ &\implies z_j'(t) = \lambda_j z_j(t),\end{aligned}$$

with $\|\mathbf{x}(t)\|^2 = \mathbf{x}(t)^*\mathbf{x}(t) = \mathbf{x}(t)^*\mathbf{U}\mathbf{U}^*\mathbf{x}(t) = \|\mathbf{z}(t)\|^2$ for all t .

The Perils of Oblique Eigenvectors

Now suppose we only have a diagonalization, $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$:

$$\begin{aligned}\mathbf{x}'(t) = \mathbf{A}\mathbf{x}(t) &\implies \mathbf{V}^{-1}\mathbf{x}'(t) = \mathbf{V}^{-1}\mathbf{A}\mathbf{V}\mathbf{V}^{-1}\mathbf{x}(t) \\ &\implies \mathbf{z}'(t) = \mathbf{\Lambda}\mathbf{z}(t) \\ &\implies z_j'(t) = \lambda_j z_j(t),\end{aligned}$$

with $\|\mathbf{x}(t)\| \neq \|\mathbf{z}(t)\|$ in general.

The Perils of Oblique Eigenvectors

Now suppose we only have a diagonalization, $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$:

$$\begin{aligned}\mathbf{x}'(t) = \mathbf{A}\mathbf{x}(t) &\implies \mathbf{V}^{-1}\mathbf{x}'(t) = \mathbf{V}^{-1}\mathbf{A}\mathbf{V}\mathbf{V}^{-1}\mathbf{x}(t) \\ &\implies \mathbf{z}'(t) = \mathbf{\Lambda}\mathbf{z}(t) \\ &\implies z_j'(t) = \lambda_j z_j(t),\end{aligned}$$

with $\|\mathbf{x}(t)\| \neq \|\mathbf{z}(t)\|$ in general.

The exact solution is easy:

$$\mathbf{x}(t) = \mathbf{V}\mathbf{z}(t) = \sum_{k=1}^n e^{t\lambda_k} z_k(0) \mathbf{v}_k.$$

Suppose $\|\mathbf{x}(0)\| = 1$. The coefficients $z_k(0)$ might still be quite large:

$$\mathbf{x}(0) = \mathbf{V}\mathbf{z}(0) = \sum_{k=0}^n z_k(0) \mathbf{v}_k.$$

The “cancellation” that gives $\mathbf{x}(0)$ is washed out for $t > 0$ by the $e^{t\lambda_k}$ terms.

Oblique Eigenvectors: Example

Example

$$\begin{bmatrix} x_1'(t) \\ x_2'(t) \end{bmatrix} = \begin{bmatrix} -1/2 & 500 \\ 0 & -5 \end{bmatrix} \begin{bmatrix} x_1(t) \\ x_2(t) \end{bmatrix}, \quad \begin{bmatrix} x_1(0) \\ x_2(0) \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}.$$

Eigenvalues and eigenvectors:

$$\lambda_1 = -1/2, \quad \mathbf{v}_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad \lambda_2 = -5, \quad \mathbf{v}_2 = \begin{bmatrix} 1 \\ -.009 \end{bmatrix}$$


Initial condition:

$$\mathbf{x}(0) = \begin{bmatrix} 1 \\ 1 \end{bmatrix} = \frac{1009}{9}\mathbf{v}_1 - \frac{1000}{9}\mathbf{v}_2.$$

Exact solution:

$$\mathbf{x}(t) = \frac{1009}{9}e^{\lambda_1 t}\mathbf{v}_1 - \frac{1000}{9}e^{\lambda_2 t}\mathbf{v}_2.$$


Oblique Eigenvectors: Example



Note the different scales of the horizontal and vertical axes.

Oblique Eigenvectors Can Lead to Transient Growth

Transient growth in the solution: a consequence of nonnormality.



NINETEEN DUBIOUS WAYS TO COMPUTE THE EXPONENTIAL OF A MATRIX*

CLEVE MOLER[†] AND CHARLES VAN LOAN[‡]




FIG. 1. The “hump”.


A Classic Paper on the Matrix Exponential

SIAM REVIEW
Vol. 20, No. 4, October 1978

© Society for Industrial and Applied Mathematics
0036-1445/78/2004-0031\$01.00/0

NINETEEN DUBIOUS WAYS TO COMPUTE THE EXPONENTIAL OF A MATRIX*

CLEVE MOLER† AND CHARLES VAN LOAN‡




Much More on Tuesday

FIG. 1. The “hump”.

Nonnormality in Iterative Linear Algebra

Nonnormality can complicate the convergence of iterative eigensolvers.




Wolfgang Kerner

'Large-scale complex eigenvalue problems'
J. Comp. Phys. 85 (1989) 1–85.

Nonnormality in Iterative Linear Algebra

Nonnormality can complicate the convergence of iterative eigensolvers.



Much More on Thursday

Wolfgang Kerner

'Large-scale complex eigenvalue problems'
J. Comp. Phys. 85 (1989) 1–85.

Tools for Measuring Nonnormality

Given a matrix, we would like some effective way to measure whether we should be concerned about the effects of nonnormality.

Tools for Measuring Nonnormality

Given a matrix, we would like some effective way to measure whether we should be concerned about the effects of nonnormality.

First, we might seek a scalar measure of normality.

Any definition of normality leads to such a gauge.

Tools for Measuring Nonnormality

Given a matrix, we would like some effective way to measure whether we should be concerned about the effects of nonnormality.

First, we might seek a scalar measure of normality.

Any definition of normality leads to such a gauge.

► $\| \mathbf{A}^* \mathbf{A} - \mathbf{A} \mathbf{A}^* \|$

Tools for Measuring Nonnormality

Given a matrix, we would like some effective way to measure whether we should be concerned about the effects of nonnormality.

First, we might seek a scalar measure of normality.

Any definition of normality leads to such a gauge.

- ▶ $\|\mathbf{A}^* \mathbf{A} - \mathbf{A} \mathbf{A}^*\|$

- ▶ $\min_{\mathbf{Z} \text{ normal}} \|\mathbf{A} - \mathbf{Z}\|$

(See work on computing the nearest normal matrix by Gabriel, Ruhe 1987.)

Tools for Measuring Nonnormality

Given a matrix, we would like some effective way to measure whether we should be concerned about the effects of nonnormality.

First, we might seek a scalar measure of normality.

Any definition of normality leads to such a gauge.

- ▶ $\|\mathbf{A}^* \mathbf{A} - \mathbf{A} \mathbf{A}^*\|$

- ▶ $\min_{\mathbf{Z} \text{ normal}} \|\mathbf{A} - \mathbf{Z}\|$

(See work on computing the nearest normal matrix by Gabriel, Ruhe 1987.)

- ▶ Henrici's departure from normality:

$$\text{dep}_2(\mathbf{A}) = \min_{\substack{\mathbf{A}=\mathbf{U}(\mathbf{D}+\mathbf{N})\mathbf{U}^* \\ \text{Schur factorization}}} \|\mathbf{N}\|.$$

No minimization is needed in the Frobenius norm:

$$\text{dep}_F(\mathbf{A}) = \min_{\substack{\mathbf{A}=\mathbf{U}(\mathbf{D}+\mathbf{N})\mathbf{U}^* \\ \text{Schur factorization}}} \|\mathbf{N}\|_F = \sqrt{\|\mathbf{A}\|_F^2 - \sum_{j=1}^n |\lambda_j|^2}.$$

Tools for Measuring Nonnormality

Given a matrix, we would like some effective way to measure whether we should be concerned about the effects of nonnormality.

First, we might seek a scalar measure of normality.

Any definition of normality leads to such a gauge.

- ▶ $\|\mathbf{A}^* \mathbf{A} - \mathbf{A} \mathbf{A}^*\|$

- ▶ $\min_{\mathbf{Z} \text{ normal}} \|\mathbf{A} - \mathbf{Z}\|$

(See work on computing the nearest normal matrix by Gabriel, Ruhe 1987.)

- ▶ Henrici's departure from normality:

$$\text{dep}_2(\mathbf{A}) = \min_{\substack{\mathbf{A}=\mathbf{U}(\mathbf{D}+\mathbf{N})\mathbf{U}^* \\ \text{Schur factorization}}} \|\mathbf{N}\|.$$

No minimization is needed in the Frobenius norm:

$$\text{dep}_F(\mathbf{A}) = \min_{\substack{\mathbf{A}=\mathbf{U}(\mathbf{D}+\mathbf{N})\mathbf{U}^* \\ \text{Schur factorization}}} \|\mathbf{N}\|_F = \sqrt{\|\mathbf{A}\|_F^2 - \sum_{j=1}^n |\lambda_j|^2}.$$

These are related by equivalence constants [Elsner, Paardekooper, 1987].

Tools for Measuring Nonnormality

Given a matrix, we would like some effective way to measure whether we should be concerned about the effects of nonnormality.

First, we might seek a scalar measure of normality.

Any definition of normality leads to such a gauge.

- ▶ $\|\mathbf{A}^* \mathbf{A} - \mathbf{A} \mathbf{A}^*\|$

- ▶ $\min_{\mathbf{Z} \text{ normal}} \|\mathbf{A} - \mathbf{Z}\|$

(See work on computing the nearest normal matrix by Gabriel, Ruhe 1987.)

- ▶ Henrici's departure from normality:

$$\text{dep}_2(\mathbf{A}) = \min_{\substack{\mathbf{A} = \mathbf{U}(\mathbf{D} + \mathbf{N})\mathbf{U}^* \\ \text{Schur factorization}}} \|\mathbf{N}\|.$$

No minimization is needed in the Frobenius norm:

$$\text{dep}_F(\mathbf{A}) = \min_{\substack{\mathbf{A} = \mathbf{U}(\mathbf{D} + \mathbf{N})\mathbf{U}^* \\ \text{Schur factorization}}} \|\mathbf{N}\|_F = \sqrt{\|\mathbf{A}\|_F^2 - \sum_{j=1}^n |\lambda_j|^2}.$$

These are related by equivalence constants [Elsner, Paardekooper, 1987].

None of these measures is of much use in practice.

Tools for Measuring Nonnormality

If $\mathbf{A}$ is diagonalizable,

$$\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1},$$

one can characterize nonnormality by $\kappa(\mathbf{V}) := \|\mathbf{V}\|\|\mathbf{V}^{-1}\| \geq 1$.

- ▶ This quantity depends on the choice of eigenvectors; scaling each column of $\mathbf{V}$ to be a unit vector gets within $\sqrt{n}$ of the optimal value, if the eigenvalues are distinct [van der Sluis 1969].
- ▶ For normal matrices, one can take $\mathbf{V}$ unitary, so $\kappa(\mathbf{V}) = 1$.
- ▶ If $\kappa(\mathbf{V})$ is not much more than 1 for some diagonalization, then the effects of nonnormality will be minimal.

Tools for Measuring Nonnormality

If $\mathbf{A}$ is diagonalizable,

$$\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1},$$

one can characterize nonnormality by $\kappa(\mathbf{V}) := \|\mathbf{V}\|\|\mathbf{V}^{-1}\| \geq 1$.

- ▶ This quantity depends on the choice of eigenvectors; scaling each column of $\mathbf{V}$ to be a unit vector gets within $\sqrt{n}$ of the optimal value, if the eigenvalues are distinct [van der Sluis 1969].
- ▶ For normal matrices, one can take $\mathbf{V}$ unitary, so $\kappa(\mathbf{V}) = 1$.
- ▶ If $\kappa(\mathbf{V})$ is not much more than 1 for some diagonalization, then the effects of nonnormality will be minimal.

Example (Bound for Continuous Systems)

$$\begin{aligned}\|\mathbf{x}(t)\| = \|\mathbf{e}^{t\mathbf{A}}\mathbf{x}(0)\| &\leq \|\mathbf{e}^{t\mathbf{A}}\|\|\mathbf{x}(0)\| \\ &\leq \|\mathbf{V}\mathbf{e}^{t\mathbf{\Lambda}}\mathbf{V}^{-1}\|\|\mathbf{x}(0)\| \\ &\leq \kappa(\mathbf{V}) \max_{\lambda \in \sigma(\mathbf{A})} |\mathbf{e}^{t\lambda}| \|\mathbf{x}(0)\|.\end{aligned}$$

1(c) Numerical range (field of values)

Rayleigh Quotients

Another approach: identify a set in the complex plane to replace the spectrum. This dates to the early 20th century literature in functional analysis, e.g., the *numerical range*, Von Neumann's *spectral sets*, and *sectorial operators*.

Definition (Rayleigh Quotient)

The *Rayleigh quotient* of $\mathbf{A} \in \mathbb{C}^{n \times n}$ with respect to nonzero $\mathbf{x} \in \mathbb{C}^n$ is

$$\frac{\mathbf{x}^* \mathbf{A} \mathbf{x}}{\mathbf{x}^* \mathbf{x}}.$$

For a Hermitian matrix $\mathbf{A}$ with eigenvalues $\lambda_1 \leq \dots \leq \lambda_n$,

$$\lambda_1 \leq \frac{\mathbf{x}^* \mathbf{A} \mathbf{x}}{\mathbf{x}^* \mathbf{x}} \leq \lambda_n.$$

More generally, Rayleigh quotients are often used as eigenvalue estimates, since if $(\lambda, \mathbf{x})$ is an eigenpair, then

$$\frac{\mathbf{x}^* \mathbf{A} \mathbf{x}}{\mathbf{x}^* \mathbf{x}} = \lambda.$$

Numerical Range (Field of Values)

So, consider looking beyond $\sigma(\mathbf{A})$ to the set of all Rayleigh quotients.

Definition (Numerical Range, a.k.a. Field of Values)

The *numerical range* of a matrix is the set

$$W(\mathbf{A}) = \{\mathbf{x}^* \mathbf{A} \mathbf{x} : \|\mathbf{x}\| = 1\}.$$

Numerical Range (Field of Values)

So, consider looking beyond $\sigma(\mathbf{A})$ to the set of all Rayleigh quotients.

Definition (Numerical Range, a.k.a. Field of Values)

The *numerical range* of a matrix is the set


$$W(\mathbf{A}) = \{\mathbf{x}^* \mathbf{A} \mathbf{x} : \|\mathbf{x}\| = 1\}.$$

- ▶ $\sigma(\mathbf{A}) \subset W(\mathbf{A})$ [Proof: take $\mathbf{x}$ to be an eigenvector.]
- ▶ $W(\mathbf{A})$ is a closed, convex subset of $\mathbf{C}$.
- ▶ If $\mathbf{U}$ is unitary, then $W(\mathbf{U}^* \mathbf{A} \mathbf{U}) = W(\mathbf{A})$.
- ▶ If $\mathbf{A}$ is normal, then $W(\mathbf{A})$ is the convex hull of $\sigma(\mathbf{A})$.
- ▶ Unlike $\sigma(\mathbf{A})$, the numerical range is robust to perturbations:


$$W(\mathbf{A} + \mathbf{E}) \subseteq W(\mathbf{A}) + \{z \in \mathbf{C} : |z| \leq \|\mathbf{E}\|\}.$$

A Gallery of Numerical Ranges


Eigenvalues and the numerical range, for four different 15×15 matrices:




normal



random



Grcar




Jordan

We will describe computation of $W(\mathbf{A})$ later this morning.

The Numerical Range Can Contain Points Far from the Spectrum


Boeing 737 example, revisited: numerical range of the 55×55 stable matrix.



We seek some middle ground between $\sigma(\mathbf{A})$ and $W(\mathbf{A})$.

The Numerical Range Can Contain Points Far from the Spectrum

Boeing 737 example, revisited: numerical range of the 55×55 stable matrix.



We seek some middle ground between $\sigma(\mathbf{A})$ and $W(\mathbf{A})$.

1(d) Pseudospectra

Pseudospectra

The numerical range $W(\mathbf{A})$ and the spectrum $\sigma(\mathbf{A})$ both have limitations. Here we shall explore another option that, loosely speaking, interpolates between $\sigma(\mathbf{A})$ and $W(\mathbf{A})$.

Pseudospectra

The numerical range $W(\mathbf{A})$ and the spectrum $\sigma(\mathbf{A})$ both have limitations. Here we shall explore another option that, loosely speaking, interpolates between $\sigma(\mathbf{A})$ and $W(\mathbf{A})$.

Example

Compute eigenvalues of three *similar* 100×100 matrices using MATLAB's eig.

$$\begin{bmatrix} 0 & 1 & & \\ 1 & 0 & \ddots & \\ & \ddots & \ddots & 1 \\ & & 1 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 0 & 1/2 & & \\ 2 & 0 & \ddots & \\ & \ddots & \ddots & 1/2 \\ & & 2 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 0 & 1/3 & & \\ 3 & 0 & \ddots & \\ & \ddots & \ddots & 1/3 \\ & & 3 & 0 \end{bmatrix}$$

Pseudospectra

The numerical range $W(\mathbf{A})$ and the spectrum $\sigma(\mathbf{A})$ both have limitations. Here we shall explore another option that, loosely speaking, interpolates between $\sigma(\mathbf{A})$ and $W(\mathbf{A})$.

Example

Compute eigenvalues of three *similar* 100×100 matrices using MATLAB's eig.


$$\begin{bmatrix} 0 & 1 & & & \\ 1 & 0 & \ddots & & \\ & \ddots & \ddots & \ddots & \\ & & \ddots & \ddots & 1 \\ & & & 1 & 0 \end{bmatrix}$$



$$\begin{bmatrix} 0 & 1/2 & & & \\ 2 & 0 & \ddots & & \\ & \ddots & \ddots & \ddots & \\ & & \ddots & \ddots & 1/2 \\ & & & 2 & 0 \end{bmatrix}$$



$$\begin{bmatrix} 0 & 1/3 & & & \\ 3 & 0 & \ddots & & \\ & \ddots & \ddots & \ddots & \\ & & \ddots & \ddots & 1/3 \\ & & & 3 & 0 \end{bmatrix}$$



Pseudospectra

Definition (ε -pseudospectrum)

For any $\varepsilon > 0$, the ε -*pseudospectrum* of $\mathbf{A}$, denoted $\sigma_\varepsilon(\mathbf{A})$, is the set

$$\sigma_\varepsilon(\mathbf{A}) = \{z \in \mathbf{C} : z \in \sigma(\mathbf{A} + \mathbf{E}) \text{ for some } \mathbf{E} \in \mathbf{C}^{n \times n} \text{ with } \|\mathbf{E}\| < \varepsilon\}.$$

Pseudospectra


Definition (ε -pseudospectrum)

For any $\varepsilon > 0$, the ε -*pseudospectrum* of $\mathbf{A}$, denoted $\sigma_\varepsilon(\mathbf{A})$, is the set

$$\sigma_\varepsilon(\mathbf{A}) = \{z \in \mathbf{C} : z \in \sigma(\mathbf{A} + \mathbf{E}) \text{ for some } \mathbf{E} \in \mathbf{C}^{n \times n} \text{ with } \|\mathbf{E}\| < \varepsilon\}.$$

We can estimate $\sigma_\varepsilon(\mathbf{A})$ by conducting experiments with random perturbations.

For the 20×20 version of $\mathbf{A} = \text{tridiag}(2, 0, 1/2)$, 50 trials each:



$$\varepsilon = 10^{-2}$$



$$\varepsilon = 10^{-4}$$



$$\varepsilon = 10^{-6}$$

Distance to Singularity

There is a fundamental connection between the distance of $\mathbf{A}$ to singularity and the norm of the inverse.

- ▶ Suppose $\|\mathbf{A}^{-1}\| = 1/\varepsilon$.
- ▶ There exists some unit vector $\mathbf{w} \in \mathbf{C}^n$ where the norm is attained:

$$\|\mathbf{A}^{-1}\| = \|\mathbf{A}^{-1}\mathbf{w}\| = \frac{1}{\varepsilon}.$$

Distance to Singularity

There is a fundamental connection between the distance of $\mathbf{A}$ to singularity and the norm of the inverse.

- ▶ Suppose $\|\mathbf{A}^{-1}\| = 1/\varepsilon$.
- ▶ There exists some unit vector $\mathbf{w} \in \mathbb{C}^n$ where the norm is attained:

$$\|\mathbf{A}^{-1}\| = \|\mathbf{A}^{-1}\mathbf{w}\| = \frac{1}{\varepsilon}.$$

- ▶ Define $\hat{\mathbf{v}} := \mathbf{A}^{-1}\mathbf{w}$.
- ▶ Define $\mathbf{v} := \hat{\mathbf{v}}/\|\hat{\mathbf{v}}\|$, a unit vector.
- ▶ $\|\mathbf{A}\mathbf{v}\| = \frac{\|\mathbf{w}\|}{\|\hat{\mathbf{v}}\|} = \varepsilon$, so $\mathbf{A}$ is “nearly singular.”

Distance to Singularity

There is a fundamental connection between the distance of $\mathbf{A}$ to singularity and the norm of the inverse.

- ▶ Suppose $\|\mathbf{A}^{-1}\| = 1/\varepsilon$.
- ▶ There exists some unit vector $\mathbf{w} \in \mathbb{C}^n$ where the norm is attained:

$$\|\mathbf{A}^{-1}\| = \|\mathbf{A}^{-1}\mathbf{w}\| = \frac{1}{\varepsilon}.$$

- ▶ Define $\hat{\mathbf{v}} := \mathbf{A}^{-1}\mathbf{w}$.
- ▶ Define $\mathbf{v} := \hat{\mathbf{v}}/\|\hat{\mathbf{v}}\|$, a unit vector.
- ▶ $\|\mathbf{A}\mathbf{v}\| = \frac{\|\mathbf{w}\|}{\|\hat{\mathbf{v}}\|} = \varepsilon$, so $\mathbf{A}$ is “nearly singular.”
- ▶ Define $\mathbf{E} := -\mathbf{A}\mathbf{v}\mathbf{v}^* \in \mathbb{C}^{n \times n}$.
- ▶ Now $(\mathbf{A} + \mathbf{E})\mathbf{v} = \mathbf{A}\mathbf{v} - \mathbf{A}\mathbf{v}\mathbf{v}^*\mathbf{v} = \mathbf{0}$, so $\mathbf{A} + \mathbf{E}$ is singular.
- ▶ The distance of $\mathbf{A}$ to singularity is $\|\mathbf{E}\| = \|\mathbf{A}\mathbf{v}\| = \varepsilon$.

Distance to Singularity

There is a fundamental connection between the distance of $\mathbf{A}$ to singularity and the norm of the inverse.

- ▶ Suppose $\|\mathbf{A}^{-1}\| = 1/\varepsilon$.
- ▶ There exists some unit vector $\mathbf{w} \in \mathbf{C}^n$ where the norm is attained:

$$\|\mathbf{A}^{-1}\| = \|\mathbf{A}^{-1}\mathbf{w}\| = \frac{1}{\varepsilon}.$$

- ▶ Define $\hat{\mathbf{v}} := \mathbf{A}^{-1}\mathbf{w}$.
- ▶ Define $\mathbf{v} := \hat{\mathbf{v}}/\|\hat{\mathbf{v}}\|$, a unit vector.
- ▶ $\|\mathbf{A}\mathbf{v}\| = \frac{\|\mathbf{w}\|}{\|\hat{\mathbf{v}}\|} = \varepsilon$, so $\mathbf{A}$ is “nearly singular.”
- ▶ Define $\mathbf{E} := -\mathbf{A}\mathbf{v}\mathbf{v}^* \in \mathbf{C}^{n \times n}$.
- ▶ Now $(\mathbf{A} + \mathbf{E})\mathbf{v} = \mathbf{A}\mathbf{v} - \mathbf{A}\mathbf{v}\mathbf{v}^*\mathbf{v} = \mathbf{0}$, so $\mathbf{A} + \mathbf{E}$ is singular.
- ▶ The distance of $\mathbf{A}$ to singularity is $\|\mathbf{E}\| = \|\mathbf{A}\mathbf{v}\| = \varepsilon$.

Norms of inverses are closely related to eigenvalue perturbations.

Equivalent Definitions of the Pseudospectrum

Theorem

The following three definitions of the ε -pseudospectrum are equivalent:

1. $\sigma_\varepsilon(\mathbf{A}) = \{z \in \mathbb{C} : z \in \sigma(\mathbf{A} + \mathbf{E}) \text{ for some } \mathbf{E} \in \mathbb{C}^{n \times n} \text{ with } \|\mathbf{E}\| < \varepsilon\};$
2. $\sigma_\varepsilon(\mathbf{A}) = \{z \in \mathbb{C} : \|(z - \mathbf{A})^{-1}\| > 1/\varepsilon\};$
3. $\sigma_\varepsilon(\mathbf{A}) = \{z \in \mathbb{C} : \|\mathbf{A}\mathbf{v} - z\mathbf{v}\| < \varepsilon \text{ for some unit vector } \mathbf{v} \in \mathbb{C}^n\}.$

Equivalent Definitions of the Pseudospectrum

Theorem

The following three definitions of the ε -pseudospectrum are equivalent:

1. $\sigma_\varepsilon(\mathbf{A}) = \{z \in \mathbb{C} : z \in \sigma(\mathbf{A} + \mathbf{E}) \text{ for some } \mathbf{E} \in \mathbb{C}^{n \times n} \text{ with } \|\mathbf{E}\| < \varepsilon\};$
2. $\sigma_\varepsilon(\mathbf{A}) = \{z \in \mathbb{C} : \|(z - \mathbf{A})^{-1}\| > 1/\varepsilon\};$
3. $\sigma_\varepsilon(\mathbf{A}) = \{z \in \mathbb{C} : \|\mathbf{A}\mathbf{v} - z\mathbf{v}\| < \varepsilon \text{ for some unit vector } \mathbf{v} \in \mathbb{C}^n\}.$

Proof. (1) $\implies$ (2)

If $z \in \sigma(\mathbf{A} + \mathbf{E})$ for some $\mathbf{E}$ with $\|\mathbf{E}\| < \varepsilon$, there exists a unit vector $\mathbf{v}$ such that $(\mathbf{A} + \mathbf{E})\mathbf{v} = z\mathbf{v}$. Rearrange to obtain

$$\mathbf{v} = (z - \mathbf{A})^{-1}\mathbf{E}\mathbf{v}.$$

Take norms:

$$\|\mathbf{v}\| = \|(z - \mathbf{A})^{-1}\mathbf{E}\mathbf{v}\| \leq \|(z - \mathbf{A})^{-1}\| \|\mathbf{E}\| \|\mathbf{v}\| < \varepsilon \|(z - \mathbf{A})^{-1}\| \|\mathbf{v}\|.$$

Hence $\|(z - \mathbf{A})^{-1}\| > 1/\varepsilon$.



Equivalent Definitions of the Pseudospectrum

Theorem

The following three definitions of the ε -pseudospectrum are equivalent:

1. $\sigma_\varepsilon(\mathbf{A}) = \{z \in \mathbb{C} : z \in \sigma(\mathbf{A} + \mathbf{E}) \text{ for some } \mathbf{E} \in \mathbb{C}^{n \times n} \text{ with } \|\mathbf{E}\| < \varepsilon\};$
2. $\sigma_\varepsilon(\mathbf{A}) = \{z \in \mathbb{C} : \|(z - \mathbf{A})^{-1}\| > 1/\varepsilon\};$
3. $\sigma_\varepsilon(\mathbf{A}) = \{z \in \mathbb{C} : \|\mathbf{A}\mathbf{v} - z\mathbf{v}\| < \varepsilon \text{ for some unit vector } \mathbf{v} \in \mathbb{C}^n\}.$

Proof. (2) $\implies$ (3)

If $\|(z - \mathbf{A})^{-1}\| > 1/\varepsilon$, there exists a unit vector $\mathbf{w}$ with $\|(z - \mathbf{A})^{-1}\mathbf{w}\| > 1/\varepsilon$. Define $\hat{\mathbf{v}} := (z - \mathbf{A})^{-1}\mathbf{w}$, so that $1/\|\hat{\mathbf{v}}\| < \varepsilon$, and

$$\frac{\|(z - \mathbf{A})\hat{\mathbf{v}}\|}{\|\hat{\mathbf{v}}\|} = \frac{\|\mathbf{w}\|}{\|\hat{\mathbf{v}}\|} < \varepsilon.$$

Hence we have found a unit vector $\mathbf{v} := \hat{\mathbf{v}}/\|\hat{\mathbf{v}}\|$ for which $\|\mathbf{A}\mathbf{v} - z\mathbf{v}\| < \varepsilon$. □

Equivalent Definitions of the Pseudospectrum

Theorem

The following three definitions of the ε -pseudospectrum are equivalent:

1. $\sigma_\varepsilon(\mathbf{A}) = \{z \in \mathbb{C} : z \in \sigma(\mathbf{A} + \mathbf{E}) \text{ for some } \mathbf{E} \in \mathbb{C}^{n \times n} \text{ with } \|\mathbf{E}\| < \varepsilon\};$
2. $\sigma_\varepsilon(\mathbf{A}) = \{z \in \mathbb{C} : \|(z - \mathbf{A})^{-1}\| > 1/\varepsilon\};$
3. $\sigma_\varepsilon(\mathbf{A}) = \{z \in \mathbb{C} : \|\mathbf{A}\mathbf{v} - z\mathbf{v}\| < \varepsilon \text{ for some unit vector } \mathbf{v} \in \mathbb{C}^n\}.$

Proof. (3) $\implies$ (1)

Given a unit vector $\mathbf{v}$ such that $\|\mathbf{A}\mathbf{v} - z\mathbf{v}\| < \varepsilon$, define $\mathbf{r} := \mathbf{A}\mathbf{v} - z\mathbf{v}$.

Now set $\mathbf{E} := -\mathbf{r}\mathbf{v}^*$, so that

$$(\mathbf{A} + \mathbf{E})\mathbf{v} = (\mathbf{A} - \mathbf{r}\mathbf{v}^*)\mathbf{v} = \mathbf{A}\mathbf{v} - \mathbf{r} = z\mathbf{v}.$$

Hence $z \in \sigma(\mathbf{A} + \mathbf{E})$.



Equivalent Definitions of the Pseudospectrum

Theorem

The following three definitions of the ε -pseudospectrum are equivalent:


1. $\sigma_\varepsilon(\mathbf{A}) = \{z \in \mathbb{C} : z \in \sigma(\mathbf{A} + \mathbf{E}) \text{ for some } \mathbf{E} \in \mathbb{C}^{n \times n} \text{ with } \|\mathbf{E}\| < \varepsilon\};$
2. $\sigma_\varepsilon(\mathbf{A}) = \{z \in \mathbb{C} : \|(z - \mathbf{A})^{-1}\| > 1/\varepsilon\};$
3. $\sigma_\varepsilon(\mathbf{A}) = \{z \in \mathbb{C} : \|\mathbf{A}\mathbf{v} - z\mathbf{v}\| < \varepsilon \text{ for some unit vector } \mathbf{v} \in \mathbb{C}^n\}.$

These different definitions are useful in different contexts:


1. interpreting numerically computed eigenvalues;
2. analyzing matrix behavior/functions of matrices;
computing pseudospectra on a grid in $\mathbb{C}$;
3. proving bounds on a particular $\sigma_\varepsilon(\mathbf{A})$.

A Gallery of Numerical Ranges


Eigenvalues and the numerical range, for four different 15×15 matrices:




normal



random




Grcar




Jordan

A Gallery of Pseudospectra


Eigenvalues and ε -pseudospectra for four different 15×15 matrices, for $\varepsilon = 10^0, 10^{-5}, 10^{-1}$:




normal



random



Grcar




Jordan

History of Pseudospectra

Invented at least four times, independently:

Jim Varah (Stanford) in his 1967 PhD thesis
Eigenvalue computations, inverse iteration



Henry Landau (Bell Labs) in a 1975 paper
Integral equations in laser theory

S. K. Godunov and colleagues (Novosibirsk), 1980s.
Eigenvalue computations, discretizations of PDEs



Nick Trefethen (MIT) starting in 1990
Discretization of PDEs, iterative linear algebra

AN INSTABILITY PHENOMENON IN SPECTRAL METHODS*

LLOYD N. TREFETHEN† AND MANFRED R. TRUMMER‡




FIG. 4. Eigenvalues of the Legendre differentiation matrix D_N^L with $N=28$, computed in various precisions ϵ . The dashes indicate the line $\operatorname{Re} \lambda = \frac{1}{2} \log_e \epsilon$.

History of Pseudospectra

Early adopters include Wilkinson (1986), Demmel (1987), Chatelin (1990s).

A Counterexample for Two Conjectures about Stability

JAMES W. DEMMEL

First computer plot of pseudospectra (1987)






Fig. 3. Contour plot of $\log_{10}(\sigma_{\min}(A - \lambda I))$ in the λ -plane.

Demmel's Pseudospectra

$$\mathbf{A} = \begin{bmatrix} -1 & -100 & -10000 \\ & -1 & -100 \\ & & -1 \end{bmatrix}$$



Find the largest ϵ value for which $\sigma_\epsilon(\mathbf{A})$ is contained in the left half plane.

The figure on the right ($\epsilon = 10^{-3}$) shows that pseudospectra can have “holes”, where the norm of the resolvent has a local minima.

Properties of Pseudospectra


$\mathbf{A}$ is normal $\iff \sigma_\varepsilon(\mathbf{A})$ is the union of open ε -balls about each eigenvalue:

$$\mathbf{A} \text{ normal} \implies \sigma_\varepsilon(\mathbf{A}) = \bigcup_j \lambda_j + \Delta_\varepsilon$$

$$\mathbf{A} \text{ nonnormal} \implies \sigma_\varepsilon(\mathbf{A}) \supset \bigcup_j \lambda_j + \Delta_\varepsilon$$

A circulant (hence normal) matrix:

$$\mathbf{A} = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \end{bmatrix}$$



Properties of Pseudospectra

$\mathbf{A}$ is normal $\iff \sigma_\varepsilon(\mathbf{A})$ is the union of open ε -balls about each eigenvalue:


$$\mathbf{A} \text{ normal} \implies \sigma_\varepsilon(\mathbf{A}) = \bigcup_j \lambda_j + \Delta_\varepsilon$$

$$\mathbf{A} \text{ nonnormal} \implies \sigma_\varepsilon(\mathbf{A}) \supset \bigcup_j \lambda_j + \Delta_\varepsilon$$

A circulant (hence normal) matrix:

$$\mathbf{A} = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \end{bmatrix}$$

$$\|(z - \mathbf{A})^{-1}\|$$



Normal Matrices: matching distance

An easy misinterpretation: "A size ε perturbation to a normal matrix can move the eigenvalues by no more than ε ."

This holds for Hermitian matrices, but not all normal matrices.

An example constructed by Gerd Krause (see Bhatia, *Matrix Analysis*, 1997):

$$\mathbf{A} = \text{diag}(1, (4 + 5\sqrt{-3})/13, (-1 + 2\sqrt{-3})/13).$$

Now construct the (unitary) Householder reflector

$$\mathbf{U} = \mathbf{I} - 2\mathbf{v}\mathbf{v}^*$$

for $\mathbf{v} = [\sqrt{5/8}, 1/2, \sqrt{1/8}]^*$, and define $\mathbf{E}$ via


$$\mathbf{A} + \mathbf{E} = -\mathbf{U}^* \mathbf{A} \mathbf{U}.$$

By construction $\mathbf{A} + \mathbf{E}$ is normal and $\sigma(\mathbf{A} + \mathbf{E}) = -\sigma(\mathbf{A})$.

Normal Matrices: matching distance

$$\mathbf{A} = \text{diag}(1, (4 + 5\sqrt{-3})/13, (-1 + 2\sqrt{-3})/13).$$


$$\mathbf{A} + \mathbf{E} = -\mathbf{U}^* \mathbf{A} \mathbf{U}.$$



Normal Matrices: matching distance

$$\mathbf{A} = \text{diag}(1, (4 + 5\sqrt{-3})/13, (-1 + 2\sqrt{-3})/13).$$

$$\mathbf{A} + \mathbf{E} = -\mathbf{U}^* \mathbf{A} \mathbf{U}.$$



Properties of Pseudospectra

Theorem (Facts about Pseudospectra)

For all $\varepsilon > 0$,

- ▶ $\sigma_\varepsilon(\mathbf{A})$ is an open, finite set that contains the spectrum.
- ▶ $\sigma_\varepsilon(\mathbf{A})$ is stable to perturbations: $\sigma_\varepsilon(\mathbf{A} + \mathbf{E}) \subseteq \sigma_{\varepsilon + \|\mathbf{E}\|}(\mathbf{A})$.
- ▶ If $\mathbf{U}$ is unitary, $\sigma_\varepsilon(\mathbf{UAU}^*) = \sigma_\varepsilon(\mathbf{A})$.
- ▶ For $\mathbf{V}$ invertible, $\sigma_{\varepsilon/\kappa(\mathbf{V})}(\mathbf{VAV}^{-1}) \subseteq \sigma_\varepsilon(\mathbf{A}) \subseteq \sigma_{\varepsilon\kappa(\mathbf{V})}(\mathbf{VAV}^{-1})$.
- ▶ $\sigma_\varepsilon(\mathbf{A} + \alpha) = \alpha + \sigma_\varepsilon(\mathbf{A})$.
- ▶ $\sigma_{|\gamma|\varepsilon}(\gamma\mathbf{A}) = \gamma\sigma_\varepsilon(\mathbf{A})$.
- ▶ $\sigma_\varepsilon\left(\begin{bmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{0} & \mathbf{B} \end{bmatrix}\right) = \sigma_\varepsilon(\mathbf{A}) \cup \sigma_\varepsilon(\mathbf{B})$.

Relationship to $\kappa(\mathbf{V})$ and $W(\mathbf{A})$

Let $\Delta_r := \{z \in \mathbb{C} : |z| < r\}$ denote the open disk of radius $r > 0$.

Theorem (Bauer–Fike, 1963)

Let $\mathbf{A}$ be diagonalizable, $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$. Then for all $\varepsilon > 0$,

$$\sigma_\varepsilon(\mathbf{A}) \subseteq \sigma(\mathbf{A}) + \Delta_{\varepsilon\kappa(\mathbf{V})}.$$

If $\kappa(\mathbf{V})$ is small, then $\sigma_\varepsilon(\mathbf{A})$ cannot contain points far from $\sigma(\mathbf{A})$.

Relationship to $\kappa(\mathbf{V})$ and $W(\mathbf{A})$

Let $\Delta_r := \{z \in \mathbb{C} : |z| < r\}$ denote the open disk of radius $r > 0$.

Theorem (Bauer–Fike, 1963)

Let $\mathbf{A}$ be diagonalizable, $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$. Then for all $\varepsilon > 0$,

$$\sigma_\varepsilon(\mathbf{A}) \subseteq \sigma(\mathbf{A}) + \Delta_{\varepsilon\kappa(\mathbf{V})}.$$

If $\kappa(\mathbf{V})$ is small, then $\sigma_\varepsilon(\mathbf{A})$ cannot contain points far from $\sigma(\mathbf{A})$.

Theorem (Stone, 1932)

For any $\mathbf{A}$,

$$\sigma_\varepsilon(\mathbf{A}) \subseteq W(\mathbf{A}) + \Delta_\varepsilon.$$

The pseudospectrum $\sigma_\varepsilon(\mathbf{A})$ cannot be bigger than $W(\mathbf{A})$ in an interesting way.

Pole Placement Example

Problem (Pole Placement Example of Mehrmann and Xu, 1996)

Given $\mathbf{A} = \text{diag}(1, 2, \dots, N)$ and $\mathbf{b} = [1, 1, \dots, 1]^T$, find $\mathbf{f} \in \mathbb{C}^n$ such that

$$\sigma(\mathbf{A} - \mathbf{b}\mathbf{f}^*) = \{-1, -2, \dots, -N\}.$$

One can show that $\mathbf{f} = \mathbf{G}^{-*}\mathbf{e}$, where $\mathbf{e} = [1, 1, \dots, 1]^T$, where

$$\mathbf{G}_{:,k} = (\mathbf{A} - \lambda_k)^{-1}\mathbf{b} = (\mathbf{A} + k)^{-1}\mathbf{b}.$$

This gives

$$\mathbf{A} - \mathbf{b}\mathbf{f}^* = \mathbf{G} \begin{bmatrix} -1 & & \\ & \ddots & \\ & & -n \end{bmatrix} \mathbf{G}^{-1},$$

with

$$\mathbf{G}_{j,k} = \frac{1}{j+k}.$$

Pole Placement Example: Numerics

In exact arithmetic, $\sigma(\mathbf{A} - \mathbf{b}\mathbf{f}^*) = \{-1, -2, \dots, -N\}$.

All entries in $\mathbf{A} - \mathbf{b}\mathbf{f}^*$ are integers.

(To ensure this, we compute $\mathbf{f}$ *symbolically*.)


For example, when $N = 8$,

$$\mathbf{A} - \mathbf{b}\mathbf{f}^* = \begin{bmatrix} 73 & -2520 & 27720 & -138600 & 360360 & -504504 & 360360 & -102960 \\ 72 & -2518 & 27720 & -138600 & 360360 & -504504 & 360360 & -102960 \\ 72 & -2520 & 27723 & -138600 & 360360 & -504504 & 360360 & -102960 \\ 72 & -2520 & 27720 & -138596 & 360360 & -504504 & 360360 & -102960 \\ 72 & -2520 & 27720 & -138600 & 360365 & -504504 & 360360 & -102960 \\ 72 & -2520 & 27720 & -138600 & 360360 & -504498 & 360360 & -102960 \\ 72 & -2520 & 27720 & -138600 & 360360 & -504504 & 360367 & -102960 \\ 72 & -2520 & 27720 & -138600 & 360360 & -504504 & 360360 & -102952 \end{bmatrix}.$$

Pole Placement Example: Numerics

In exact arithmetic, $\sigma(\mathbf{A} - \mathbf{b}\mathbf{f}^*) = \{-1, -2, \dots, -N\}$.


Computed eigenvalues of $\mathbf{A} - \mathbf{b}\mathbf{f}^*$ (matrix is exact: all integer entries).



Pole Placement Example: Numerics

In exact arithmetic, $\sigma(\mathbf{A} - \mathbf{b}\mathbf{f}^*) = \{-1, -2, \dots, -N\}$.


Computed eigenvalues of $\mathbf{A} - \mathbf{b}\mathbf{f}^*$ (matrix is exact: all integer entries).



Pole Placement Example: Numerics

In exact arithmetic, $\sigma(\mathbf{A} - \mathbf{b}\mathbf{f}^*) = \{-1, -2, \dots, -N\}$.


Computed eigenvalues of $\mathbf{A} - \mathbf{b}\mathbf{f}^*$ (matrix is exact: all integer entries).



Pole Placement Example: Numerics

In exact arithmetic, $\sigma(\mathbf{A} - \mathbf{b}\mathbf{f}^*) = \{-1, -2, \dots, -N\}$.


Computed eigenvalues of $\mathbf{A} - \mathbf{b}\mathbf{f}^*$ (matrix is exact: all integer entries).



Pole Placement Example: Numerics

In exact arithmetic, $\sigma(\mathbf{A} - \mathbf{b}\mathbf{f}^*) = \{-1, -2, \dots, -N\}$.


Computed eigenvalues of $\mathbf{A} - \mathbf{b}\mathbf{f}^*$ (matrix is exact: all integer entries).



Pole Placement Example: Numerics

In exact arithmetic, $\sigma(\mathbf{A} - \mathbf{b}\mathbf{f}^*) = \{-1, -2, \dots, -N\}$.


Computed eigenvalues of $\mathbf{A} - \mathbf{b}\mathbf{f}^*$ (matrix is exact: all integer entries).



Pole Placement Example: Numerics

In exact arithmetic, $\sigma(\mathbf{A} - \mathbf{b}\mathbf{f}^*) = \{-1, -2, \dots, -N\}$.


Computed eigenvalues of $\mathbf{A} - \mathbf{b}\mathbf{f}^*$ (matrix is exact: all integer entries).



Pole Placement Example: Numerics

In exact arithmetic, $\sigma(\mathbf{A} - \mathbf{b}\mathbf{f}^*) = \{-1, -2, \dots, -N\}$.


Computed eigenvalues of $\mathbf{A} - \mathbf{b}\mathbf{f}^*$ (matrix is exact: all integer entries).



Pole Placement Example: Numerics

In exact arithmetic, $\sigma(\mathbf{A} - \mathbf{b}\mathbf{f}^*) = \{-1, -2, \dots, -N\}$.


Computed eigenvalues of $\mathbf{A} - \mathbf{b}\mathbf{f}^*$ (matrix is exact: all integer entries).



Pole Placement Example: Numerics

In exact arithmetic, $\sigma(\mathbf{A} - \mathbf{b}\mathbf{f}^*) = \{-1, -2, \dots, -N\}$.


Computed eigenvalues of $\mathbf{A} - \mathbf{b}\mathbf{f}^*$ (matrix is exact: all integer entries).



Pole Placement Example: Numerics


In exact arithmetic, $\sigma(\mathbf{A} - \mathbf{b}\mathbf{f}^*) = \{-1, -2, \dots, -N\}$.

Computed eigenvalues of $\mathbf{A} - \mathbf{b}\mathbf{f}^*$ (matrix is exact: all integer entries).



Pole Placement Example: Pseudospectra


Computed pseudospectra of $\mathbf{A} - \mathbf{b}\mathbf{f}^*$.



Computed eigenvalues should be accurate to roughly $\varepsilon_{\text{mach}} \|\mathbf{A} - \mathbf{b}\mathbf{f}^*\|$.
For example, when $N = 11$, $\|\mathbf{A} - \mathbf{b}\mathbf{f}^*\| = 5.26 \times 10^8$.

Pole Placement Example: Pseudospectra


Computed pseudospectra of $\mathbf{A} - \mathbf{b}\mathbf{f}^*$.



Computed eigenvalues should be accurate to roughly $\varepsilon_{\text{mach}} \|\mathbf{A} - \mathbf{b}\mathbf{f}^*\|$.
For example, when $N = 11$, $\|\mathbf{A} - \mathbf{b}\mathbf{f}^*\| = 5.26 \times 10^8$.

Pole Placement Example: Pseudospectra


Computed pseudospectra of $\mathbf{A} - \mathbf{b}\mathbf{f}^*$.



Computed eigenvalues should be accurate to roughly $\varepsilon_{\text{mach}} \|\mathbf{A} - \mathbf{b}\mathbf{f}^*\|$.
For example, when $N = 11$, $\|\mathbf{A} - \mathbf{b}\mathbf{f}^*\| = 5.26 \times 10^8$.

Pole Placement Example: Pseudospectra


Computed pseudospectra of $\mathbf{A} - \mathbf{b}\mathbf{f}^*$.



Computed eigenvalues should be accurate to roughly $\varepsilon_{\text{mach}} \|\mathbf{A} - \mathbf{b}\mathbf{f}^*\|$.
For example, when $N = 11$, $\|\mathbf{A} - \mathbf{b}\mathbf{f}^*\| = 5.26 \times 10^8$.

Pole Placement Example: Pseudospectra


Computed pseudospectra of $\mathbf{A} - \mathbf{b}\mathbf{f}^*$.



Computed eigenvalues should be accurate to roughly $\varepsilon_{\text{mach}} \|\mathbf{A} - \mathbf{b}\mathbf{f}^*\|$.
For example, when $N = 11$, $\|\mathbf{A} - \mathbf{b}\mathbf{f}^*\| = 5.26 \times 10^8$.

Pole Placement Example: Pseudospectra


Computed pseudospectra of $\mathbf{A} - \mathbf{b}\mathbf{f}^*$.



Computed eigenvalues should be accurate to roughly $\varepsilon_{\text{mach}} \|\mathbf{A} - \mathbf{b}\mathbf{f}^*\|$.
For example, when $N = 11$, $\|\mathbf{A} - \mathbf{b}\mathbf{f}^*\| = 5.26 \times 10^8$.

Pole Placement Example: Pseudospectra


Computed pseudospectra of $\mathbf{A} - \mathbf{b}\mathbf{f}^*$.



Computed eigenvalues should be accurate to roughly $\varepsilon_{\text{mach}} \|\mathbf{A} - \mathbf{b}\mathbf{f}^*\|$.
For example, when $N = 11$, $\|\mathbf{A} - \mathbf{b}\mathbf{f}^*\| = 5.26 \times 10^8$.

Pole Placement Example: Pseudospectra


Computed pseudospectra of $\mathbf{A} - \mathbf{b}\mathbf{f}^*$.



Computed eigenvalues should be accurate to roughly $\varepsilon_{\text{mach}} \|\mathbf{A} - \mathbf{b}\mathbf{f}^*\|$.
For example, when $N = 11$, $\|\mathbf{A} - \mathbf{b}\mathbf{f}^*\| = 5.26 \times 10^8$.

Pole Placement Example: Pseudospectra

Computed pseudospectra of $\mathbf{A} - \mathbf{b}\mathbf{f}^*$.



Computed eigenvalues should be accurate to roughly $\varepsilon_{\text{mach}} \|\mathbf{A} - \mathbf{b}\mathbf{f}^*\|$.
For example, when $N = 11$, $\|\mathbf{A} - \mathbf{b}\mathbf{f}^*\| = 5.26 \times 10^8$.

Polynomial Zeros and Companion Matrices

MATLAB's `roots` command computes polynomial zeros by computing the eigenvalues of a companion matrix.

For example, given $p(z) = c_0 + c_1z + c_2z^2 + c_3z^3 + c_4z^4$, MATLAB builds

$$\mathbf{A} = \begin{bmatrix} -c_4/c_0 & -c_3/c_0 & -c_2/c_0 & -c_1/c_0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}.$$

whose characteristic polynomial is p .

Polynomial Zeros and Companion Matrices

MATLAB's `roots` command computes polynomial zeros by computing the eigenvalues of a companion matrix.

For example, given $p(z) = c_0 + c_1z + c_2z^2 + c_3z^3 + c_4z^4$, MATLAB builds

$$\mathbf{A} = \begin{bmatrix} -c_4/c_0 & -c_3/c_0 & -c_2/c_0 & -c_1/c_0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}.$$

whose characteristic polynomial is p .

Problem (Wilkinson's "Perfidious Polynomial")

Find the zeros of the polynomial


$$p(z) = (z - 1)(z - 2) \cdots (z - N)$$

from coefficients in the monomial basis.

MATLAB gives this as an example: `roots(poly(1:20))`.


Roots of Wilkinson's Perfidious Polynomial

Computed eigenvalues of the companion matrix.




Roots of Wilkinson's Perfidious Polynomial

Computed eigenvalues of the companion matrix.




Roots of Wilkinson's Perfidious Polynomial

Computed eigenvalues of the companion matrix.




Roots of Wilkinson's Perfidious Polynomial

Computed eigenvalues of the companion matrix.




Roots of Wilkinson's Perfidious Polynomial

Computed eigenvalues of the companion matrix.




Roots of Wilkinson's Perfidious Polynomial

Pseudospectra for $N = 25$.




Roots of Wilkinson's Perfidious Polynomial

3d plot of the resolvent norm reveals the a local minimum.




1(e) Computing $W(\mathbf{A})$ and $\sigma_\varepsilon(\mathbf{A})$


A Gallery of Numerical Ranges




normal



random




Grcar




Jordan


A Gallery of Numerical Ranges




normal



random



Grcar



Jordan

If $z \in W(\mathbf{A})$, then

$$\operatorname{Re} z = \frac{z + \bar{z}}{2} = \frac{1}{2}(\mathbf{x}^* \mathbf{A} \mathbf{x} + \mathbf{x}^* \mathbf{A}^* \mathbf{x}) = \mathbf{x}^* \left(\frac{\mathbf{A} + \mathbf{A}^*}{2} \right) \mathbf{x}.$$


Using properties of Hermitian matrices, we conclude that

$$\operatorname{Re}(W(\mathbf{A})) = \left[\lambda_{\min} \left(\frac{\mathbf{A} + \mathbf{A}^*}{2} \right), \lambda_{\max} \left(\frac{\mathbf{A} + \mathbf{A}^*}{2} \right) \right].$$

Similarly, one can determine the intersection of $W(\mathbf{A})$ with any line in $\mathbb{C}$.


Computation of the Numerical Range

This calculation yields points on the boundary of the numerical range.
Use convexity to obtain polygonal outer and inner approximations
[Johnson 1980]; Higham's `fv.m`.



Computation of the Numerical Range

This calculation yields points on the boundary of the numerical range.
Use convexity to obtain polygonal outer and inner approximations
[Johnson 1980]; Higham's `fv.m`.




Neat problem: Given $z \in W(\mathbf{A})$, find unit vector $\mathbf{x}$ such that $z = \mathbf{x}^* \mathbf{A} \mathbf{x}$
[Uhlig 2008; Carden 2009].

Computation of Pseudospectra (Dense Matrices)

Naive algorithm: $\mathcal{O}(n^3)$ per grid point


- ▶ Compute $\|(z - \mathbf{A})^{-1}\|$ using the SVD on a grid of points in $\mathbb{C}$.
SVD costs $\mathcal{O}(n^3)$ for dense $\mathbf{A}$.
- ▶ Send data to a contour plotting routine.



Computation of Pseudospectra (Dense Matrices)

Naive algorithm: $\mathcal{O}(n^3)$ per grid point

- ▶ Compute $\|(z - \mathbf{A})^{-1}\|$ using the SVD on a grid of points in $\mathbb{C}$.
SVD costs $\mathcal{O}(n^3)$ for dense $\mathbf{A}$.
- ▶ Send data to a contour plotting routine.



Computation of Pseudospectra (Dense Matrices)

Modern algorithm: $\mathcal{O}(n^3) + \mathcal{O}(n^2)$ per grid point [Lui 1997; Trefethen 1999]

- ▶ Compute a Schur triangularization, $\mathbf{A} = \mathbf{UTU}^*$.
Schur costs $\mathcal{O}(n^3)$ for dense $\mathbf{A}$.
- ▶ Recall that $\sigma_\varepsilon(\mathbf{A}) = \sigma_\varepsilon(\mathbf{T})$.

Computation of Pseudospectra (Dense Matrices)

Modern algorithm: $\mathcal{O}(n^3) + \mathcal{O}(n^2)$ per grid point [Lui 1997; Trefethen 1999]

- ▶ Compute a Schur triangularization, $\mathbf{A} = \mathbf{U}\mathbf{T}\mathbf{U}^*$.
Schur costs $\mathcal{O}(n^3)$ for dense $\mathbf{A}$.
- ▶ Recall that $\sigma_\varepsilon(\mathbf{A}) = \sigma_\varepsilon(\mathbf{T})$.
- ▶ One can efficiently solve linear systems $\mathbf{T}\mathbf{x} = \mathbf{b}$ for $\mathbf{x}$.
Each solve costs $\mathcal{O}(n^2)$ for upper triangular $\mathbf{T}$.

Computation of Pseudospectra (Dense Matrices)

Modern algorithm: $\mathcal{O}(n^3) + \mathcal{O}(n^2)$ per grid point [Lui 1997; Trefethen 1999]

- ▶ Compute a Schur triangularization, $\mathbf{A} = \mathbf{UTU}^*$.
Schur costs $\mathcal{O}(n^3)$ for dense $\mathbf{A}$.
- ▶ Recall that $\sigma_\varepsilon(\mathbf{A}) = \sigma_\varepsilon(\mathbf{T})$.
- ▶ One can efficiently solve linear systems $\mathbf{T}\mathbf{x} = \mathbf{b}$ for $\mathbf{x}$.
Each solve costs $\mathcal{O}(n^2)$ for upper triangular $\mathbf{T}$.
- ▶ To compute $\|(z - \mathbf{T})^{-1}\| = 1/s_{\min}(z - \mathbf{T}) = s_{\max}((z - \mathbf{T})^{-1})$,
find the largest eigenvalue of $(z - \mathbf{T})^{-*}(z - \mathbf{T})^{-1}$.
- ▶ To do so, apply the *Lanczos iteration* on $(z - \mathbf{T})^{-*}(z - \mathbf{T})^{-1}$
for each grid point z .

Computation of Pseudospectra (Dense Matrices)

Modern algorithm: $\mathcal{O}(n^3) + \mathcal{O}(n^2)$ per grid point [Lui 1997; Trefethen 1999]

- ▶ Compute a Schur triangularization, $\mathbf{A} = \mathbf{UTU}^*$.
Schur costs $\mathcal{O}(n^3)$ for dense $\mathbf{A}$.
- ▶ Recall that $\sigma_\varepsilon(\mathbf{A}) = \sigma_\varepsilon(\mathbf{T})$.
- ▶ One can efficiently solve linear systems $\mathbf{T}\mathbf{x} = \mathbf{b}$ for $\mathbf{x}$.
Each solve costs $\mathcal{O}(n^2)$ for upper triangular $\mathbf{T}$.
- ▶ To compute $\|(z - \mathbf{T})^{-1}\| = 1/s_{\min}(z - \mathbf{T}) = s_{\max}((z - \mathbf{T})^{-1})$,
find the largest eigenvalue of $(z - \mathbf{T})^{-*}(z - \mathbf{T})^{-1}$.
- ▶ To do so, apply the *Lanczos iteration* on $(z - \mathbf{T})^{-*}(z - \mathbf{T})^{-1}$
for each grid point z .
- ▶ At each step, one must compute $(z - \mathbf{T})^{-*}(z - \mathbf{T})^{-1}\mathbf{b}$:
This requires a lower-triangular solve with $(z - \mathbf{T})^*$ ($\mathcal{O}(n^2)$)
and an upper triangular solve with $(z - \mathbf{T})$ ($\mathcal{O}(n^2)$).
One observes a constant number of iterations per grid point.

Computation of Pseudospectra (Dense Matrices)

Modern algorithm: $\mathcal{O}(n^3) + \mathcal{O}(n^2)$ per grid point [Lui 1997; Trefethen 1999]

- ▶ Compute a Schur triangularization, $\mathbf{A} = \mathbf{UTU}^*$.
Schur costs $\mathcal{O}(n^3)$ for dense $\mathbf{A}$.
- ▶ Recall that $\sigma_\varepsilon(\mathbf{A}) = \sigma_\varepsilon(\mathbf{T})$.
- ▶ One can efficiently solve linear systems $\mathbf{T}\mathbf{x} = \mathbf{b}$ for $\mathbf{x}$.
Each solve costs $\mathcal{O}(n^2)$ for upper triangular $\mathbf{T}$.
- ▶ To compute $\|(z - \mathbf{T})^{-1}\| = 1/s_{\min}(z - \mathbf{T}) = s_{\max}((z - \mathbf{T})^{-1})$,
find the largest eigenvalue of $(z - \mathbf{T})^{-*}(z - \mathbf{T})^{-1}$.
- ▶ To do so, apply the *Lanczos iteration* on $(z - \mathbf{T})^{-*}(z - \mathbf{T})^{-1}$
for each grid point z .
- ▶ At each step, one must compute $(z - \mathbf{T})^{-*}(z - \mathbf{T})^{-1}\mathbf{b}$:
This requires a lower-triangular solve with $(z - \mathbf{T})^*$ ($\mathcal{O}(n^2)$)
and an upper triangular solve with $(z - \mathbf{T})$ ($\mathcal{O}(n^2)$).
One observes a constant number of iterations per grid point.
- ▶ Total cost: $\mathcal{O}(n^3) + \mathcal{O}(n^2)$ per grid point.

Computation of Pseudospectra (Sparse Matrices)

Large-scale problems [Toh and Trefethen 1996; Wright and Trefethen 2001]




Key idea:

- ▶ Find $\mathbf{V} \in \mathbb{C}^{n \times k}$ with orthonormal columns, $\mathbf{V}^* \mathbf{A} \mathbf{V}$, for $k \ll n$.
- ▶ The “generalized Rayleigh quotient” $\mathbf{V}^* \mathbf{A} \mathbf{V} \in \mathbb{C}^{k \times k}$.
- ▶ Approximate $\sigma_\varepsilon(\mathbf{V}^* \mathbf{A} \mathbf{V}) \approx \sigma_\varepsilon(\mathbf{A})$.

In general, $\sigma(\mathbf{V}^* \mathbf{A} \mathbf{V}) \not\subseteq \sigma(\mathbf{A})$, so for some $\varepsilon > 0$,

$$\sigma_\varepsilon(\mathbf{V}^* \mathbf{A} \mathbf{V}) \not\subseteq \sigma_\varepsilon(\mathbf{A}).$$

Depending on the choice of $\mathbf{V}$, the approximation $\sigma_\varepsilon(\mathbf{V}^* \mathbf{A} \mathbf{V})$ might give a rough general impression of $\sigma_\varepsilon(\mathbf{A})$, or it might give a rather accurate approximation in one interesting region of $\sigma_\varepsilon(\mathbf{A})$.



Computation of Pseudospectra (Sparse Matrices)

First important choice for $\mathbf{V}_k$ [Toh and Trefethen 1996]:

- Projection onto Krylov subspaces (Arnoldi factorization)

$$\text{Ran}(\mathbf{V}_k) = \text{span}\{\mathbf{x}, \mathbf{A}\mathbf{x}, \dots, \mathbf{A}^{k-1}\mathbf{x}\}.$$

The **Arnoldi process** generates orthonormal bases for Krylov subspaces:

$$\mathbf{A}\mathbf{V}_k = \mathbf{V}_{k+1}\tilde{\mathbf{H}}_k,$$

where $\tilde{\mathbf{H}}_k \in \mathbf{C}^{(k+1) \times k}$ is upper Hessenberg.

Due to the upper Hessenberg structure, we have

$$s_{\min}(z - \tilde{\mathbf{H}}_k) \geq s_{\min}(z - \mathbf{A}),$$

so one can get a *bound*

$$\sigma_{\varepsilon}(\tilde{\mathbf{H}}_k) \subseteq \sigma_{\varepsilon}(\mathbf{A}),$$

where, for a rectangular matrix $\mathbf{R}$, we have

$$\sigma_{\varepsilon}(\mathbf{R}) := \{z \in \mathbf{C} : s_{\min}(z - \mathbf{R}) < \varepsilon\}.$$

Computation of Pseudospectra (Sparse Matrices)

Second important choice for $\mathbf{V}_k$ [Wright and Trefethen 2001]:

- Projection onto an invariant subspace (eigenspace)

$$\mathbf{V}_k = [\mathbf{v}_1, \dots, \mathbf{v}_k]$$

Suppose $\mathbf{A}\mathbf{V}_k = \mathbf{V}_k\mathbf{X}$ for some $\mathbf{X} \in \mathbb{C}^{k \times k}$.

Then $\text{Ran}(\mathbf{V}_k)$ is an *invariant subspace* of $\mathbf{A}$.

If $\mathbf{V} = [\mathbf{V}_k \ \widehat{\mathbf{V}}_k]$ is a unitary matrix, $\mathbf{V}^*\mathbf{V} = \mathbf{I}$, then

$$\mathbf{V}^*\mathbf{A}\mathbf{V} = \begin{bmatrix} \mathbf{V}_k^*\mathbf{A}\mathbf{V}_k & \widehat{\mathbf{V}}_k^*\mathbf{A}\mathbf{V}_k \\ \mathbf{0} & \widehat{\mathbf{V}}_k^*\mathbf{A}\widehat{\mathbf{V}}_k \end{bmatrix}.$$

Thus $\sigma_\varepsilon(\mathbf{V}_k^*\mathbf{A}\mathbf{V}_k) \subseteq \sigma_\varepsilon(\mathbf{A})$.

Compute an invariant subspace corresponding to eigenvalues of physical interest (e.g., using ARPACK).

Computation of Pseudospectra

Alternative: [Brühl 1996; Bekas and Gallopoulos, ...]

Curve tracing: follow level sets of $\|(z - \mathbf{A})^{-1}\|$.

- ▶ Given a point $z = x + iy \in \mathbf{C}$, suppose $\|z - \mathbf{A}\|^{-1} = 1/\varepsilon$.
- ▶ Suppose the smallest singular value s of $z - \mathbf{A}$ is simple, with singular vectors $\mathbf{u}$ and $\mathbf{v}$:

$$(z - \mathbf{A})\mathbf{v} = s\mathbf{u}.$$

- ▶ Brühl uses a result of Sun (1988) to obtain

$$\frac{\partial s}{\partial x} = \operatorname{Re}(\mathbf{u}^* \mathbf{v}), \quad \frac{\partial s}{\partial y} = \operatorname{Im}(\mathbf{v}^* \mathbf{u}).$$

- ▶ Use these derivatives to follow the boundary $\partial\sigma_\varepsilon(\mathbf{A})$.

Alternative: [Brühl 1996; Bekas and Gallopoulos, ...]

Curve tracing: follow level sets of $\|(z - \mathbf{A})^{-1}\|$.

- ▶ Given a point $z = x + iy \in \mathbf{C}$, suppose $\|z - \mathbf{A}\|^{-1} = 1/\varepsilon$.
- ▶ Suppose the smallest singular value s of $z - \mathbf{A}$ is simple, with singular vectors $\mathbf{u}$ and $\mathbf{v}$:

$$(z - \mathbf{A})\mathbf{v} = s\mathbf{u}.$$

- ▶ Brühl uses a result of Sun (1988) to obtain


$$\frac{\partial s}{\partial x} = \operatorname{Re}(\mathbf{u}^* \mathbf{v}), \quad \frac{\partial s}{\partial y} = \operatorname{Im}(\mathbf{v}^* \mathbf{u}).$$

- ▶ Use these derivatives to follow the boundary $\partial\sigma_\varepsilon(\mathbf{A})$.

This seems like an elegant alternative to the grid-based method, but:

- ▶ It only gives $\sigma_\varepsilon(\mathbf{A})$ for one value of ε .
- ▶ One must beware of cusps, holes, disconnected components of $\sigma_\varepsilon(\mathbf{A})$.

EigTool: Software for Pseudospectra Computation



EigTool: Thomas Wright, 2002

<http://www.cs.ox.ac.uk/pseudospectra/eigtool>