Math 6610: Analysis of Numerical Methods, I
Eigenvalues and eigenvectors

Department of Mathematics, University of Utah

Fall 2025

Resources: Trefethen and Bau 1997, Lecture 24
Süli and Mayers 2003, Section 5.1
Salgado and Wise 2022, Sections 1.3, 8.3

Eigenvalues and eigenvectors

Given $\boldsymbol{A} \in \mathbb{C}^{n \times n}$, $(\lambda, \boldsymbol{v}) \in \mathbb{C} \times (\mathbb{C}^n \backslash \{0\})$ is an eigenvalue-eigenvector pair if

$$\boldsymbol{A}\boldsymbol{v} = \lambda \boldsymbol{v}.$$

Recall: it doesn't matter what value(s) λ takes on, but $\boldsymbol{v}$ *cannot* be $\mathbf{0}$.

Eigenvalues and eigenvectors

Given $\boldsymbol{A} \in \mathbb{C}^{n \times n}$, $(\lambda, \boldsymbol{v}) \in \mathbb{C} \times (\mathbb{C}^n \backslash \{0\})$ is an eigenvalue-eigenvector pair if

$$\boldsymbol{A}\boldsymbol{v} = \lambda \boldsymbol{v}.$$

Recall: it doesn't matter what value(s) λ takes on, but $\boldsymbol{v}$ *cannot* be $\mathbf{0}$.

Some additional terminology/properties:

- The collection of all eigenvalues of $\boldsymbol{A}$ is $\lambda(\boldsymbol{A}) \subset \mathbb{C}$, its *spectrum*.
- Even if $\boldsymbol{A} \in \mathbb{R}^{n \times n}$, $\lambda(\boldsymbol{A})$ can contain complex values, and eigenvectors can be complex-valued.
- On paper, we typically identify the spectrum by computing roots of the *characteristic polynomial*, $p_{\boldsymbol{A}}(\lambda) = \det(\lambda \boldsymbol{I} - \boldsymbol{A})$. (This turns out to be a *terrible* algorithmic strategy.)

Eigenpair properties

There are many properties of eigenvalues and eigenvectors of $\boldsymbol{A} \in \mathbb{C}^{n \times n}$:

Eigenpair properties

There are many properties of eigenvalues and eigenvectors of $\boldsymbol{A} \in \mathbb{C}^{n \times n}$:

- All square matrices have exactly n eigenvalues, with some possibly repeated.
- All square matrices have at least 1 eigenvector.

Eigenpair properties

There are many properties of eigenvalues and eigenvectors of $\boldsymbol{A} \in \mathbb{C}^{n \times n}$:

- All square matrices have exactly n eigenvalues, with some possibly repeated.
- All square matrices have at least 1 eigenvector.
- Eigenvectors from *distinct* eigenvalues are linearly independent.
- Every *distinct* eigenvalue has at least 1 eigenvector.

Eigenpair properties

There are many properties of eigenvalues and eigenvectors of $\boldsymbol{A} \in \mathbb{C}^{n \times n}$:

- All square matrices have exactly n eigenvalues, with some possibly repeated.
- All square matrices have at least 1 eigenvector.
- Eigenvectors from *distinct* eigenvalues are linearly independent.
- Every *distinct* eigenvalue has at least 1 eigenvector.
- The collection of all eigenvectors associated to an eigenvalue λ is a subspace, and is frequently called an *eigenspace* E_λ.
- Eigenspaces are invariant subspaces of $\boldsymbol{A}$, i.e., $\boldsymbol{A}E_\lambda \subseteq E_\lambda$.

Eigenpair properties

There are many properties of eigenvalues and eigenvectors of $\boldsymbol{A} \in \mathbb{C}^{n \times n}$:

- All square matrices have exactly n eigenvalues, with some possibly repeated.
- All square matrices have at least 1 eigenvector.
- Eigenvectors from *distinct* eigenvalues are linearly independent.
- Every *distinct* eigenvalue has at least 1 eigenvector.
- The collection of all eigenvectors associated to an eigenvalue λ is a subspace, and is frequently called an *eigenspace* E_λ.
- Eigenspaces are invariant subspaces of $\boldsymbol{A}$, i.e., $\boldsymbol{A}E_\lambda \subseteq E_\lambda$.
- The number of times an eigenvalue is repeated a_λ is its *algebraic multiplicity*
- The *geometric multiplicity* g_λ of an eigenvalue λ is $\dim E_\lambda$.
- *Simple* eigenvalues λ have $g_\lambda = a_\lambda = 1$.

Eigenpair properties

There are many properties of eigenvalues and eigenvectors of $\boldsymbol{A} \in \mathbb{C}^{n \times n}$:

- All square matrices have exactly n eigenvalues, with some possibly repeated.
- All square matrices have at least 1 eigenvector.
- Eigenvectors from *distinct* eigenvalues are linearly independent.
- Every *distinct* eigenvalue has at least 1 eigenvector.
- The collection of all eigenvectors associated to an eigenvalue λ is a subspace, and is frequently called an *eigenspace* E_λ.
- Eigenspaces are invariant subspaces of $\boldsymbol{A}$, i.e., $\boldsymbol{A}E_\lambda \subseteq E_\lambda$.
- The number of times an eigenvalue is repeated a_λ is its *algebraic multiplicity*
- The *geometric multiplicity* g_λ of an eigenvalue λ is $\dim E_\lambda$.
- *Simple* eigenvalues λ have $g_\lambda = a_\lambda = 1$.
- In general we always have $g_\lambda \leqslant a_\lambda$, so that, $1 \leqslant \sum_{\lambda \in \lambda(\boldsymbol{A})} g_\lambda \leqslant \sum_{\lambda \in \lambda(\boldsymbol{A})} a_\lambda = n$.

Eigenpair properties

There are many properties of eigenvalues and eigenvectors of $\boldsymbol{A} \in \mathbb{C}^{n \times n}$:

- All square matrices have exactly n eigenvalues, with some possibly repeated.
- All square matrices have at least 1 eigenvector.
- Eigenvectors from *distinct* eigenvalues are linearly independent.
- Every *distinct* eigenvalue has at least 1 eigenvector.
- The collection of all eigenvectors associated to an eigenvalue λ is a subspace, and is frequently called an *eigenspace* E_λ.
- Eigenspaces are invariant subspaces of $\boldsymbol{A}$, i.e., $\boldsymbol{A}E_\lambda \subseteq E_\lambda$.
- The number of times an eigenvalue is repeated a_λ is its *algebraic multiplicity*
- The *geometric multiplicity* g_λ of an eigenvalue λ is $\dim E_\lambda$.
- *Simple* eigenvalues λ have $g_\lambda = a_\lambda = 1$.
- In general we always have $g_\lambda \leqslant a_\lambda$, so that, $1 \leqslant \sum_{\lambda \in \lambda(\boldsymbol{A})} g_\lambda \leqslant \sum_{\lambda \in \lambda(\boldsymbol{A})} a_\lambda = n$.
- Eigenvalues λ with $g_\lambda < a_\lambda$ are *defective*
- Any $\boldsymbol{A}$ with a defective eigenvalue is *defective*.

A digression: similarity transforms

Two square matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ are *similar* if $\exists$ an invertible $\boldsymbol{S}$ such that

$$\boldsymbol{B} = \boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{S}.$$

(The map $\boldsymbol{A} \mapsto \boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{S}$ is a similarity transform.)

A digression: similarity transforms

Two square matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ are *similar* if $\exists$ an invertible $\boldsymbol{S}$ such that

$$\boldsymbol{B} = \boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{S}.$$

(The map $\boldsymbol{A} \mapsto \boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{S}$ is a similarity transform.)

Suppose $\boldsymbol{A}$ is not defective: then with linearly independent $\boldsymbol{v}_i$, $i \in [n]$ we have the relations,

$$\boldsymbol{A}\boldsymbol{v}_i = \lambda_i \boldsymbol{v}_i.$$

This is equivalent to,

$$\boldsymbol{A}\boldsymbol{V} = \boldsymbol{\Lambda}\boldsymbol{V}, \qquad \boldsymbol{V} = \begin{pmatrix} | & & | \\ \boldsymbol{v}_1 & \cdots & \boldsymbol{v}_n \\ | & & | \end{pmatrix}, \qquad \boldsymbol{\Lambda} = \mathrm{diag}(\lambda_1, \ldots, \lambda_n).$$

A digression: similarity transforms

Two square matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ are *similar* if $\exists$ an invertible $\boldsymbol{S}$ such that

$$\boldsymbol{B} = \boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{S}.$$

(The map $\boldsymbol{A} \mapsto \boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{S}$ is a similarity transform.)

Suppose $\boldsymbol{A}$ is not defective: then with linearly independent $\boldsymbol{v}_i$, $i \in [n]$ we have the relations,

$$\boldsymbol{A}\boldsymbol{v}_i = \lambda_i \boldsymbol{v}_i.$$

This is equivalent to,

$$\boldsymbol{A}\boldsymbol{V} = \boldsymbol{\Lambda}\boldsymbol{V}, \qquad \boldsymbol{V} = \begin{pmatrix} | & & | \\ \boldsymbol{v}_1 & \cdots & \boldsymbol{v}_n \\ | & & | \end{pmatrix}, \qquad \boldsymbol{\Lambda} = \mathrm{diag}(\lambda_1, \ldots, \lambda_n).$$

And since $\boldsymbol{V}$ is full-rank, then

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

i.e., $\boldsymbol{A}$ is similar to a diagonal matrix containing its spectrum.

Matrix diagonalizability

This motivates a key definition and consequence.

Definition

A square matrix $\boldsymbol{A} \in \mathbb{C}^{n \times n}$ is diagonalizable if it is similar to a diagonal matrix.

Theorem

A square matrix $\boldsymbol{A}$ is diagonalizable iff it is not defective.

When $\boldsymbol{A}$ is not defective, it is diagonalizable via a matrix whose columns are comprised of its linearly independent eigenvectors.

Similarity invariances

One simple observation is that the set of eigenvalues is invariant under a(ny) similarity transform, since if $\boldsymbol{A}$ is diagonalizable with diagonal matrix $\boldsymbol{\Lambda}$, then,

$$\boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{S} = (\boldsymbol{V}\boldsymbol{S})^{-1}\,\boldsymbol{\Lambda}\,(\boldsymbol{V}\boldsymbol{S})$$

is a diagonalization of $\boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{S}$ with the same diagonal matrix $\boldsymbol{\Lambda}$.

Similarity invariances

One simple observation is that the set of eigenvalues is invariant under a(ny) similarity transform, since if $\boldsymbol{A}$ is diagonalizable with diagonal matrix $\boldsymbol{\Lambda}$, then,

$$\boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{S} = (\boldsymbol{V}\boldsymbol{S})^{-1}\,\boldsymbol{\Lambda}\,(\boldsymbol{V}\boldsymbol{S})$$

is a diagonalization of $\boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{S}$ with the same diagonal matrix $\boldsymbol{\Lambda}$.

More consequences follow, e.g., some familiar determinant and trace properties,

$$\det \boldsymbol{A} = \prod_{j=1}^{n} \lambda_j, \qquad \operatorname{tr}\boldsymbol{A} = \sum_{j=1}^{n} \lambda_j.$$

The above is actually true for any square matrix $\boldsymbol{A}$, defective or not.

Jordan normal form

Defective matrices certainly exist. The most common example is,

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}.$$

A rather nice fact is that the above example is "spiritually" the prototypical and essentially only type of defective matrix.

Defective matrices certainly exist. The most common example is,

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}.$$

A rather nice fact is that the above example is "spiritually" the prototypical and essentially only type of defective matrix.

Theorem (Jordan normal form)

Every square matrix $\boldsymbol{A} \in \mathbb{C}^{n \times n}$ is bidiagonalizable, i.e., is similar to a bidiagonal matrix.

More specifically, let $\lambda_1, \ldots, \lambda_U$ be the unique eigenvalues of $\boldsymbol{A}$, and suppose they have algebraic and geometric multiplicities a_j and g_j, $j \in [U]$, respectively. Tnen:

$$\boldsymbol{A} = \boldsymbol{V}\boldsymbol{J}\boldsymbol{V}^{-1}, \qquad \boldsymbol{V} \in \mathbb{C}^{n \times n},$$

where $\boldsymbol{V}$ is invertible, and $\boldsymbol{J}$ is bidiagonal with $\Lambda(\boldsymbol{A})$ on the diagonal, given by,

$$\boldsymbol{J} = \bigoplus_{j \in [U]} \left(\lambda_j \boldsymbol{I}_{g_j - 1} \oplus \boldsymbol{J}_j\right), \qquad \boldsymbol{J}_j = \lambda_j \boldsymbol{I}_{a_j - g_j + 1} + \boldsymbol{N}_{a_j - g_j + 1},$$

where $\boldsymbol{N}_k$ is a $k \times k$ matrix, nonzero only on its main superdiagonal that has entries all 1.

To diagonalizability and beyond

"Most" square matrices are diagonalizable.

This is (incredibly) powerful: a symmetric linear change of basis of the input and output spaces results in a diagonal linear operator.
The particular change of basis can be quite anisotropic (and non-isometric) in nature.

All square matrices are not diagonalizable. However, ...

To diagonalizability and beyond

"Most" square matrices are diagonalizable.

This is (incredibly) powerful: a symmetric linear change of basis of the input and output spaces results in a diagonal linear operator.
The particular change of basis can be quite anisotropic (and non-isometric) in nature.

All square matrices are not diagonalizable. However, ...

- ...all square matrices are bidiagonalizable (Jordan normal form)

To diagonalizability and beyond

"Most" square matrices are diagonalizable.

This is (incredibly) powerful: a symmetric linear change of basis of the input and output spaces results in a diagonal linear operator.
The particular change of basis can be quite anisotropic (and non-isometric) in nature.

All square matrices are not diagonalizable. However, ...

- ...all square matrices are bidiagonalizable (Jordan normal form)
- ...all square matrices are unitarily triangularizable (Schur decomposition)
  While triangularizability is not as clean as diagonalizability, that there are unitary transformations accomplishing this is very attractive.

To diagonalizability and beyond

"Most" square matrices are diagonalizable.

This is (incredibly) powerful: a symmetric linear change of basis of the input and output spaces results in a diagonal linear operator.
The particular change of basis can be quite anisotropic (and non-isometric) in nature.

All square matrices are not diagonalizable. However, ...

- ...all square matrices are bidiagonalizable (Jordan normal form)
- ...all square matrices are unitarily triangularizable (Schur decomposition)
  While triangularizability is not as clean as diagonalizability, that there are unitary transformations accomplishing this is very attractive.

Are there matrices that are *unitarily* diagonalizable?

A "spectral" theorem

A seemingly unrelated algebraic definition is our starting point.

Definition

A matrix $\boldsymbol{A} \in \mathbb{C}^{n \times n}$ is *Hermitian* if $\boldsymbol{A} = \boldsymbol{A}^*$.
(Hermitian matrices are also called *self-adjoint*, or *symmetric* when $\boldsymbol{A}$ is real-valued.)

A "spectral" theorem

A seemingly unrelated algebraic definition is our starting point.

Definition

A matrix $\boldsymbol{A} \in \mathbb{C}^{n \times n}$ is *Hermitian* if $\boldsymbol{A} = \boldsymbol{A}^*$.
(Hermitian matrices are also called *self-adjoint*, or *symmetric* when $\boldsymbol{A}$ is real-valued.)

Theorem (Spectral Theorem for Hermitian matrices)

If $\boldsymbol{A} \in \mathbb{C}^{n \times n}$ is Hermitian, then it is unitarily diagonalizable with real eigenvalues.
(Its spectrum is real-valued, and the similarity matrix accomplishing diagonalization is unitary.)

A “spectral” theorem

A seemingly unrelated algebraic definition is our starting point.

Definition

A matrix $\boldsymbol{A} \in \mathbb{C}^{n \times n}$ is *Hermitian* if $\boldsymbol{A} = \boldsymbol{A}^*$.
(Hermitian matrices are also called *self-adjoint*, or *symmetric* when $\boldsymbol{A}$ is real-valued.)

Theorem (Spectral Theorem for Hermitian matrices)

If $\boldsymbol{A} \in \mathbb{C}^{n \times n}$ is Hermitian, then it is unitarily diagonalizable with real eigenvalues.
(Its spectrum is real-valued, and the similarity matrix accomplishing diagonalization is unitary.)

Hermitian matrices are *very* common in applications, and the spectral theorem has numerous uses.

Application I: Geometric interpretation of Hermitian matrices

If $\boldsymbol{A} \in \mathbb{C}^{n \times n}$ is unitarily diagonalizable, then it can be written as

$$\boldsymbol{A} = \boldsymbol{U} \boldsymbol{\Lambda} \boldsymbol{U}^* = \sum_{j=1}^{n} \lambda_j \boldsymbol{u}_j \boldsymbol{u}_j^*,$$

where $\{\boldsymbol{u}_j\}_{j=1}^{n}$ are the columns of $\boldsymbol{U}$.

Application I: Geometric interpretation of Hermitian matrices

If $\boldsymbol{A} \in \mathbb{C}^{n \times n}$ is unitarily diagonalizable, then it can be written as

$$\boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{U}^* = \sum_{j=1}^{n} \lambda_j \boldsymbol{u}_j \boldsymbol{u}_j^*,$$

where $\{\boldsymbol{u}_j\}_{j=1}^{n}$ are the columns of $\boldsymbol{U}$.

I.e., Hermitian matrices (an algebraic property) have strong geometric interpretation: they are "just" diagonal matrices in a rotated/reflected orthonormal frame.

Application II: The induced 2-norm

The *spectral radius* of a matrix $\boldsymbol{A}$ is

$$\rho(\boldsymbol{A}) := \max_{j=1,\ldots,n} |\lambda_j(\boldsymbol{A})|$$

If $\boldsymbol{A}$ is Hermitian, then $\|\boldsymbol{A}\|_2 = \rho(\boldsymbol{A})$.

This is direct from the definition of the induced 2-norm.

Application III: The "$\boldsymbol{A}$ norm"

A matrix $\boldsymbol{A} \in \mathbb{C}^{n \times n}$ is Hermitian positive definite (sometimes *symmetric* positive-definite or "spd") if it's Hermitian and its (real) spectrum is strictly positive.

(Respectively, positive semi-definite if the spectrum is non-negative.)

Application III: The "$\boldsymbol{A}$ norm"

A matrix $\boldsymbol{A} \in \mathbb{C}^{n \times n}$ is Hermitian positive definite (sometimes *symmetric* positive-definite or "spd") if it's Hermitian and its (real) spectrum is strictly positive.

(Respectively, positive semi-definite if the spectrum is non-negative.)

Such matrices actually define a norm: $\|\boldsymbol{x}\|_{\boldsymbol{A}}^2 := \boldsymbol{x}^* \boldsymbol{A} \boldsymbol{x}$ is a norm.

Application IV: Matrix square roots

There is a functional calculus on spd matrices.

For example, a matrix $\boldsymbol{B}$ is the square root of a matrix $\boldsymbol{A}$ if $\boldsymbol{A} = \boldsymbol{B}^2$.

Application IV: Matrix square roots

There is a functional calculus on spd matrices.

For example, a matrix $\boldsymbol{B}$ is the square root of a matrix $\boldsymbol{A}$ if $\boldsymbol{A} = \boldsymbol{B}^2$.

Example

If $\boldsymbol{A}$ is spd, compute a matrix square root of $\boldsymbol{A}$.

Application IV: Matrix square roots

There is a functional calculus on spd matrices.

For example, a matrix $\boldsymbol{B}$ is the square root of a matrix $\boldsymbol{A}$ if $\boldsymbol{A} = \boldsymbol{B}^2$.

Example

If $\boldsymbol{A}$ is spd, compute a matrix square root of $\boldsymbol{A}$.

Theorem

If $\boldsymbol{A}$ is spd, then there is a unique spd square root $\boldsymbol{B}$ of $\boldsymbol{A}$, i.e., $\boldsymbol{B}^2 = \boldsymbol{A}$.

Application IV: Quadratic forms

Given a Hermitian matrix $\boldsymbol{A} \in \mathbb{C}^{n \times n}$, the function,

$$Q_{\boldsymbol{A}}(\boldsymbol{x}) := \boldsymbol{x}^* \boldsymbol{A} \boldsymbol{x},$$

is a **quadratic form**, i.e., a real-valued quadratic function on $\mathbb{C}^n$. The eigendecomposition of $\boldsymbol{A}$ uniquely defines the behavior of $Q_{\boldsymbol{A}}$.

Application IV: Quadratic forms

Given a Hermitian matrix $\boldsymbol{A} \in \mathbb{C}^{n \times n}$, the function,

$$Q_{\boldsymbol{A}}(\boldsymbol{x}) := \boldsymbol{x}^* \boldsymbol{A} \boldsymbol{x},$$

is a **quadratic form**, i.e., a real-valued quadratic function on $\mathbb{C}^n$. The eigendecomposition of $\boldsymbol{A}$ uniquely defines the behavior of $Q_{\boldsymbol{A}}$.

If the following eigenvectors correspond to the positive, negative, and zero eigenvalues of $\boldsymbol{A}$, respectively,

$$\left\{\boldsymbol{v}_i^+\right\}_{i \in [n^+]}, \qquad \left\{\boldsymbol{v}_i^-\right\}_{i \in [n^-]}, \qquad \left\{\boldsymbol{v}_i^0\right\}_{i \in [n^0]},$$

where $n = n^+ + n^- + n^0$. Then clearly:

$$Q_{\boldsymbol{A}}(\boldsymbol{v}_i^+) > 0, \qquad Q_{\boldsymbol{A}}(\boldsymbol{v}_i^-) < 0, \qquad Q_{\boldsymbol{A}}(\boldsymbol{v}_i^0) = 0.$$

Application IV: Quadratic forms

Given a Hermitian matrix $\boldsymbol{A} \in \mathbb{C}^{n\times n}$, the function,

$$Q_{\boldsymbol{A}}(\boldsymbol{x}) := \boldsymbol{x}^* \boldsymbol{A}\boldsymbol{x},$$

is a **quadratic form**, i.e., a real-valued quadratic function on $\mathbb{C}^n$. The eigendecomposition of $\boldsymbol{A}$ uniquely defines the behavior of $Q_{\boldsymbol{A}}$.

If the following eigenvectors correspond to the positive, negative, and zero eigenvalues of $\boldsymbol{A}$, respectively,

$$\left\{\boldsymbol{v}_i^+\right\}_{i\in[n^+]}, \qquad \left\{\boldsymbol{v}_i^-\right\}_{i\in[n^-]}, \qquad \left\{\boldsymbol{v}_i^0\right\}_{i\in[n^0]},$$

where $n = n^+ + n^- + n^0$. Then clearly:

$$Q_{\boldsymbol{A}}(\boldsymbol{v}_i^+) > 0, \qquad Q_{\boldsymbol{A}}(\boldsymbol{v}_i^-) < 0, \qquad Q_{\boldsymbol{A}}(\boldsymbol{v}_i^0) = 0.$$

Generalizing this a bit:

$$\left.\begin{array}{l} V^+ := \left\{\boldsymbol{v}_i^+\right\}_{i\in[n^+]}, \\ V^- := \left\{\boldsymbol{v}_i^-\right\}_{i\in[n^-]}, \\ V^0 := \left\{\boldsymbol{v}_i^0\right\}_{i\in[n^0]} \end{array}\right\} \Longrightarrow \left\{\begin{array}{l} Q_{\boldsymbol{A}}(\boldsymbol{x}) > 0 \text{ if } \boldsymbol{x} \in V^+\backslash\{\boldsymbol{0}\} \\ Q_{\boldsymbol{A}}(\boldsymbol{x}) < 0 \text{ if } \boldsymbol{x} \in V^-\backslash\{\boldsymbol{0}\} \\ Q_{\boldsymbol{A}}(\boldsymbol{x}) = 0 \text{ if } \boldsymbol{x} \in V^0 \end{array}\right.$$

where $\mathbb{C}^n = V^+ \oplus V^- \oplus V^0$.

Rayleigh quotients, I

A final application of Hermitian matrices is a *variational* characterization of eigenvalues. We need some buildup for this.

Let $\boldsymbol{A} \in \mathbb{C}^{n \times n}$ be a(ny) square matrix, and let $\boldsymbol{x} \in \mathbb{C}^n \backslash \{\boldsymbol{0}\}$ be a vector.

The Rayleigh Quotient (of $\boldsymbol{A}$ at $\boldsymbol{x}$) is the (complex) scalar,

$$R_{\boldsymbol{A}}(\boldsymbol{x}) := \frac{Q_{\boldsymbol{A}}(\boldsymbol{x})}{\|\boldsymbol{x}\|_2^2} = \frac{\boldsymbol{x}^* \boldsymbol{A} \boldsymbol{x}}{\boldsymbol{x}^* \boldsymbol{x}} = \frac{\langle \boldsymbol{A}\boldsymbol{x}, \boldsymbol{x} \rangle}{\langle \boldsymbol{x}, \boldsymbol{x} \rangle}, \qquad \boldsymbol{x} \neq \boldsymbol{0}$$

Ostensibly, if $(\lambda, \boldsymbol{v})$ is an eigenpair of $\boldsymbol{A}$, then $R_{\boldsymbol{A}}(\boldsymbol{v}) = \lambda$.

Rayleigh quotients, I

A final application of Hermitian matrices is a *variational* characterization of eigenvalues. We need some buildup for this.

Let $\boldsymbol{A} \in \mathbb{C}^{n \times n}$ be a(ny) square matrix, and let $\boldsymbol{x} \in \mathbb{C}^n \backslash \{\boldsymbol{0}\}$ be a vector.

The Rayleigh Quotient (of $\boldsymbol{A}$ at x) is the (complex) scalar,

$$R_{\boldsymbol{A}}(\boldsymbol{x}) := \frac{Q_{\boldsymbol{A}}(\boldsymbol{x})}{\|\boldsymbol{x}\|_2^2} = \frac{\boldsymbol{x}^* \boldsymbol{A} \boldsymbol{x}}{\boldsymbol{x}^* \boldsymbol{x}} = \frac{\langle \boldsymbol{A}\boldsymbol{x}, \boldsymbol{x} \rangle}{\langle \boldsymbol{x}, \boldsymbol{x} \rangle}, \qquad \boldsymbol{x} \neq \boldsymbol{0}$$

Ostensibly, if $(\lambda, \boldsymbol{v})$ is an eigenpair of $\boldsymbol{A}$, then $R_{\boldsymbol{A}}(\boldsymbol{v}) = \lambda$.

The numerical range of $\boldsymbol{A}$ is the set of all possible values of $R_{\boldsymbol{A}}$:

$$W(\boldsymbol{A}) := R_{\boldsymbol{A}}\left(\mathbb{C}^n \backslash \{\boldsymbol{0}\}\right).$$

One can view $W(\boldsymbol{A})$ as the image of the Rayleigh quotient over all $\mathbb{C}^n$, but also just as the image of the Rayleigh quotient over the unit sphere in $\mathbb{C}^n$.

$W(\boldsymbol{A})$ is some set in $\mathbb{C}$, regardless of the dimension n of $\boldsymbol{A}$.

Clearly we know $\lambda(\boldsymbol{A}) \subset W(\boldsymbol{A})$.

Rayleigh quotients, II

There is a rather more interesting property of the numerical range.

Theorem (Hausdorff-Toeplitz Theorem)

$W(\boldsymbol{A})$ *is a compact and convex set in* $\mathbb{C}$.

Compactness: $W(\boldsymbol{A})$ is the image of a compact set (unit sphere in $\mathbb{C}^n$) under a continuous function ($R_{\boldsymbol{A}}(\cdot)$).

Rayleigh quotients, II

There is a rather more interesting property of the numerical range.

Theorem (Hausdorff-Toeplitz Theorem)

$W(\boldsymbol{A})$ is a compact and convex set in $\mathbb{C}$.

Compactness: $W(\boldsymbol{A})$ is the image of a compact set (unit sphere in $\mathbb{C}^n$) under a continuous function ($R_{\boldsymbol{A}}(\cdot)$).

For certain classes of matrices, the Rayleigh quotient is a little more transparent.

For example, if $\boldsymbol{A}$ is Hermitian, then $R_{\boldsymbol{A}}(\boldsymbol{x}) \in \mathbb{R}$, so $W(\boldsymbol{A}) \subset \mathbb{R}$.

Rayleigh quotients, II

There is a rather more interesting property of the numerical range.

Theorem (Hausdorff-Toeplitz Theorem)

$W(\boldsymbol{A})$ is a compact and convex set in $\mathbb{C}$.

Compactness: $W(\boldsymbol{A})$ is the image of a compact set (unit sphere in $\mathbb{C}^n$) under a continuous function ($R_{\boldsymbol{A}}(\cdot)$).

For certain classes of matrices, the Rayleigh quotient is a little more transparent.

For example, if $\boldsymbol{A}$ is Hermitian, then $R_{\boldsymbol{A}}(\boldsymbol{x}) \in \mathbb{R}$, so $W(\boldsymbol{A}) \subset \mathbb{R}$.

In fact, something more precise is true

Theorem

If $\boldsymbol{A}$ is Hermitian, then

$$\lambda_{\min}(\boldsymbol{A}) \leqslant R_{\boldsymbol{A}}(\boldsymbol{x}) \leqslant \lambda_{\max}(\boldsymbol{A}), \qquad \boldsymbol{x} \in \mathbb{C}^n \backslash \{\boldsymbol{0}\}.$$

An immediate corollary: If $\boldsymbol{A}$ is Hermitian, then $W(\boldsymbol{A}) = [\lambda_{\min}(\boldsymbol{A}), \lambda_{\max}(\boldsymbol{A})]$.

Rayleigh quotients, III

Let $\boldsymbol{A} \in \mathbb{C}^{n \times n}$ be Hermitian. Consider a subspace $V \subset \mathbb{C}^n$.

The image of the V under the Rayleigh quotient, $R_{\boldsymbol{A}}(V)$, is some subset of $W(\boldsymbol{A}) \subset \mathbb{R}$.

Rayleigh quotients, III

Let $\boldsymbol{A} \in \mathbb{C}^{n \times n}$ be Hermitian. Consider a subspace $V \subset \mathbb{C}^n$.

The image of the V under the Rayleigh quotient, $R_{\boldsymbol{A}}(V)$, is some subset of $W(\boldsymbol{A}) \subset \mathbb{R}$.

- The minimum value of $R_{\boldsymbol{A}}(V)$ is $\lambda_{\min}(\boldsymbol{A})$, occuring when V contains the minimum eigenvector. What is the largest possible minimum value?

Rayleigh quotients, III

Let $\boldsymbol{A} \in \mathbb{C}^{n \times n}$ be Hermitian. Consider a subspace $V \subset \mathbb{C}^n$.

The image of the V under the Rayleigh quotient, $R_{\boldsymbol{A}}(V)$, is some subset of $W(\boldsymbol{A}) \subset \mathbb{R}$.

- The minimum value of $R_{\boldsymbol{A}}(V)$ is $\lambda_{\min}(\boldsymbol{A})$, occuring when V contains the minimum eigenvector.
  What is the largest possible minimum value?
- The maximum value of $R_{\boldsymbol{A}}(V)$ is $\lambda_{\max}(\boldsymbol{A})$.
  What is the smallest possible maximum value?

The "min-max" theorem

Theorem (Courant-Fischer-Weyl "min-max")

Let $\boldsymbol{A} \in \mathbb{C}^{n\times n}$ *be Hermitian, with eigenvalues* $\lambda_1 \leqslant \lambda_2 \leqslant \ldots \leqslant \lambda_n$. *Then for each* $1 \leqslant k \leqslant n$,

$$\lambda_k = \min_{\substack{V \subset \mathbb{C}^n \\ \dim V = k}} \max W_{\boldsymbol{A}}(V)$$

$$\lambda_k = \max_{\substack{V \subset \mathbb{C}^n \\ \dim V = n-k+1}} \min W_{\boldsymbol{A}}(V)$$

In addition, if $(\boldsymbol{u}_j)_{j=1}^n$ *are the eigenvectors associated with* $(\lambda_j)_{j=1}^n$, *then:*

- $V = \mathrm{span}\{\boldsymbol{u}_1, \ldots, \boldsymbol{u}_k\}$ *achieves the outer minimum*
- $V = \mathrm{span}\{\boldsymbol{u}_k, \ldots, \boldsymbol{u}_n\}$ *achieves the outer maximum*

Cauchy interlacing theorem

A matrix $\boldsymbol{B}$ is a **compression** of $\boldsymbol{A}$ if $\boldsymbol{B} = \boldsymbol{Q}^* \boldsymbol{A} \boldsymbol{Q}$ for some $\boldsymbol{Q} \in \mathbb{C}^{n \times r}$ with orthonormal columns.

Cauchy interlacing theorem

A matrix $\boldsymbol{B}$ is a **compression** of $\boldsymbol{A}$ if $\boldsymbol{B} = \boldsymbol{Q}^* \boldsymbol{A} \boldsymbol{Q}$ for some $\boldsymbol{Q} \in \mathbb{C}^{n \times r}$ with orthonormal columns.

Just one consequence of the min-max theorem:

Theorem (Cauchy interlacing)

Let $\boldsymbol{B} \in \mathbb{C}^{(n-1)\times(n-1)}$ be a compression of a Hermitian matrix $\boldsymbol{A} \in \mathbb{C}^{n \times n}$. If $\boldsymbol{A}$ has eigenvalues $\lambda_1 \leqslant \lambda_2 \leqslant \ldots \leqslant \lambda_n$, and $\boldsymbol{B}$ has eigenvalues $\mu_1, \ldots, \mu_{n-1}$, then

$$\lambda_j \leqslant \mu_j \leqslant \lambda_{j+1},$$

for all $j = 1, \ldots, n-1$.