DSC 140B

Representation Learning

Lecture 03 | Part 1

News

News

- Cheat sheet allowed on quizzes.
- FinAid survey on Gradescope.
  - If you’re here in person (and participate in the Live Q&A), you don’t need to do this.
- Practice problems now on `dsc140b.com`

DSC 140B
Representation Learning

Lecture 03 | Part 2

Functions of a Vector

Transformations

- A **transformation** $\vec{f}$ is a function that takes in a vector, and returns a vector *of the same dimensionality*.

- That is, $\vec{f} : \mathbb{R}^d \to \mathbb{R}^d$.

Arbitrary Transformations

- Arbitrary transformations can be quite complex.

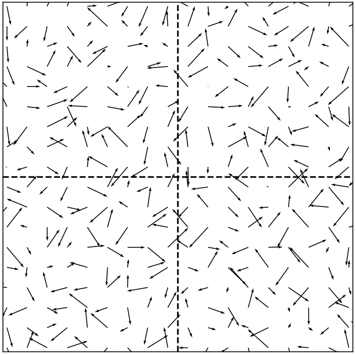

Arbitrary Transformations

- Arbitrary transformations can be quite complex.

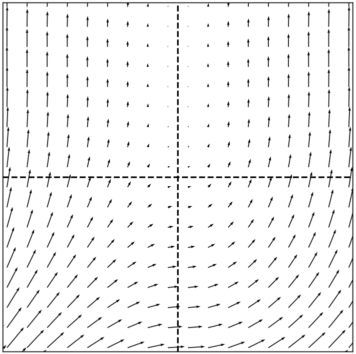

Linear Transformations

- Luckily, we often[1] work with simpler, **linear transformations**.

- A transformation f is **linear** if:

$$\vec{f}(\alpha\vec{x} + \beta\vec{y}) = \alpha\vec{f}(\vec{x}) + \beta\vec{f}(\vec{y})$$

[1]Sometimes just to make the math tractable!

Key Fact

- Linear functions are determined **entirely** by what they do on the basis vectors.
- I.e., to tell you what f does, I only need to tell you $\vec{f}(\hat{e}^{(1)})$ and $\vec{f}(\hat{e}^{(2)})$.
- This makes the math easy!

Linear Algebra

- This is the key idea behind **linear** algebra.
- Linear algebra studies the properties of **linear** transformations.
- Non-linear transformations are **so complicated** that we can say relatively little about them.

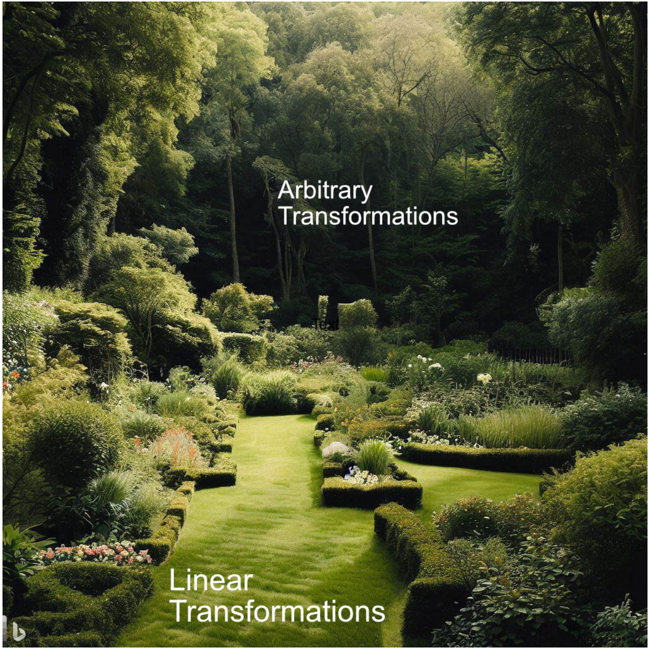
Arbitrary
Transformations
Linear
Transformations

Example Linear Transformation

- $\vec{f}(\vec{x}) = (x_1 + 3x_2, -3x_1 + 5x_2)^T$

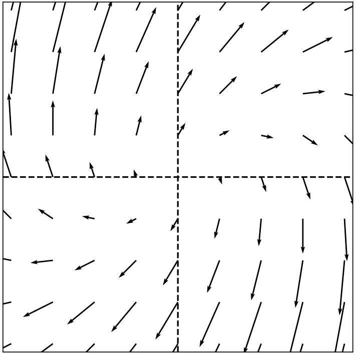

Another Example Linear Transformation

- $\vec{f}(\vec{x}) = (2x_1 - x_2, -x_1 + 3x_2)^T$

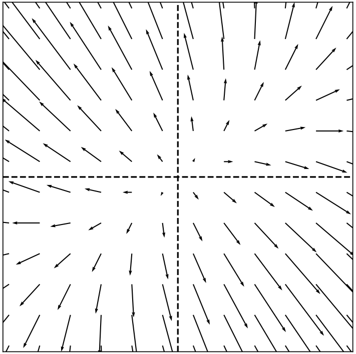

Note

- Because of linearity, along any given direction $\vec{f}$ changes only in scale.

$$\vec{f}(\lambda \hat{x}) = \lambda \vec{f}(\hat{x})$$

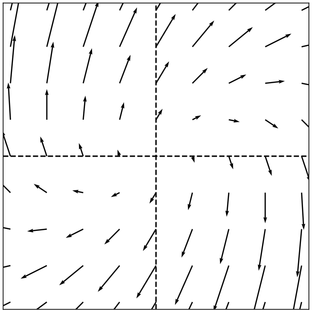

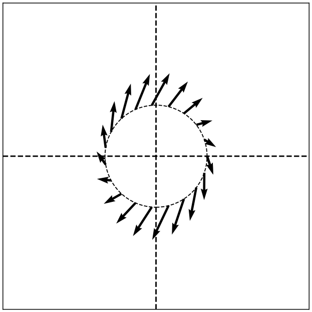

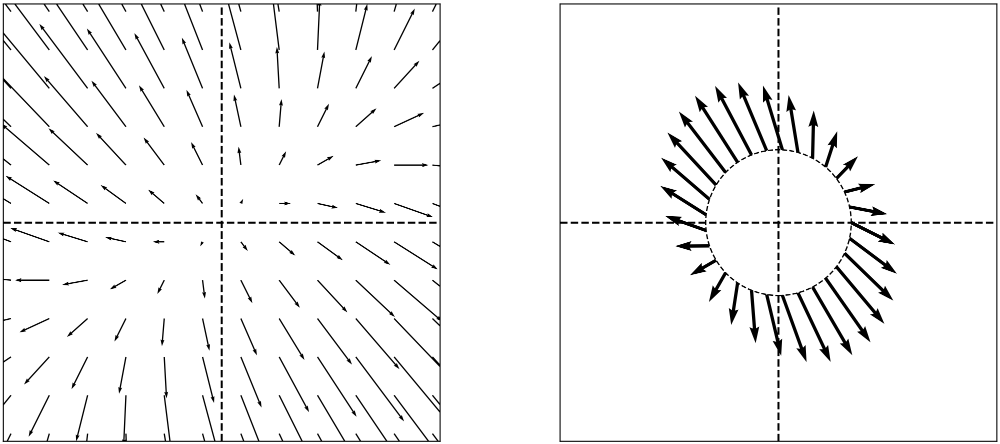

Linear Transformations and Bases

- We have been writing transformations in coordinate form. For example:

$$\begin{aligned}\vec{f}(\vec{x}) &= (x_1 + x_2, x_1 - x_2)^T \\ &= (x_1 + x_2)\hat{e}^{(1)} + (x_1 - x_2)\hat{e}^{(2)}\end{aligned}$$

- If we use a different basis, the formula for $\vec{f}$ **changes**:

$$\begin{aligned}[\vec{f}(\vec{x})]_{\mathcal{U}} &= (?, ?)^T \\ &= [?]\hat{u}^{(1)} + [?]\hat{u}^{(2)}\end{aligned}$$

Linear Transformations and Bases

- We know that if $\vec{x} = x_1 \hat{e}^{(1)} + x_2 \hat{e}^{(2)}$, then:

$$\vec{f}(\vec{x}) = (x_1 + x_2)\hat{e}^{(1)} + (x_1 - x_2)\hat{e}^{(2)}$$

- Now: if $\vec{x} = z_1 \hat{u}^{(1)} + z_2 \hat{u}^{(2)}$, what is:

$$\vec{f}(\vec{x}) = ?\hat{u}^{(1)} + ?\hat{u}^{(2)}$$

Key Fact

- If we use linearity:

$$\begin{aligned} f(\vec{x}) &= f(z_1 \hat{u}^{(1)} + z_2 \hat{u}^{(2)}) \\ &= z_1 f(\hat{u}^{(1)}) + z_2 f(\hat{u}^{(2)}) \end{aligned}$$

- **Strategy**: to write $\vec{f}$ in the $\mathcal{U}$ basis, we just need to know what $\vec{f}$ does to $\hat{u}^{(1)}$ and $\hat{u}^{(2)}$.

Example

- Let:
  - $\vec{f}(\vec{x}) = (x_1 + x_2, x_1 - x_2)^T$
  - $\hat{u}^{(1)} = \frac{1}{\sqrt{2}}(1, 1)^T$ and $\hat{u}^{(2)} = \frac{1}{\sqrt{2}}(-1, 1)^T$.
- Then:

$$\vec{f}(\hat{u}^{(1)}) = \vec{f}\left(\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}\right)^T = \left(\sqrt{2}, 0\right)^T = \sqrt{2}\hat{e}^{(1)}$$

$$\vec{f}(\hat{u}^{(2)}) = \vec{f}\left(-\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}\right)^T = \left(0, -\sqrt{2}\right)^T = -\sqrt{2}\hat{e}^{(2)}$$

- **But** we want $\vec{f}(\hat{u}^{(1)})$ and $\vec{f}(\hat{u}^{(2)})$ in terms of $\hat{u}^{(1)}$ and $\hat{u}^{(2)}$.

Example (Cont.)

- We have: $f(\hat{u}^{(1)}) = \sqrt{2}\hat{e}^{(1)}$ and $f(\hat{u}^{(2)}) = -\sqrt{2}\hat{e}^{(2)}$.
- To write $\vec{f}(\hat{u}^{(1)})$ in terms of $\hat{u}^{(1)}$ and $\hat{u}^{(2)}$, compute:

$$
\begin{aligned}
f(\hat{u}^{(1)}) &= (f(\hat{u}^{(1)}) \cdot \hat{u}^{(1)})\hat{u}^{(1)} + (f(\hat{u}^{(1)}) \cdot \hat{u}^{(2)})\hat{u}^{(2)} \\
&= \\
&=
\end{aligned}
$$

Example (Cont.)

- We have: $f(\hat{u}^{(1)}) = \sqrt{2}\hat{e}^{(1)}$ and $f(\hat{u}^{(2)}) = -\sqrt{2}\hat{e}^{(2)}$.
- To write $\vec{f}(\hat{u}^{(1)})$ in terms of $\hat{u}^{(1)}$ and $\hat{u}^{(2)}$, compute:

$$\begin{aligned} f(\hat{u}^{(1)}) &= (f(\hat{u}^{(1)}) \cdot \hat{u}^{(1)})\hat{u}^{(1)} + (f(\hat{u}^{(1)}) \cdot \hat{u}^{(2)})\hat{u}^{(2)} \\ &= \left((\sqrt{2}, 0) \cdot \frac{1}{\sqrt{2}}(1, 1)\right)\hat{u}^{(1)} + \left((\sqrt{2}, 0) \cdot \frac{1}{\sqrt{2}}(-1, 1)\right)\hat{u}^{(2)} \\ &= \end{aligned}$$

Example (Cont.)

- We have: $f(\hat{u}^{(1)}) = \sqrt{2}\hat{e}^{(1)}$ and $f(\hat{u}^{(2)}) = -\sqrt{2}\hat{e}^{(2)}$.
- To write $\vec{f}(\hat{u}^{(1)})$ in terms of $\hat{u}^{(1)}$ and $\hat{u}^{(2)}$, compute:

$$\begin{aligned} f(\hat{u}^{(1)}) &= (f(\hat{u}^{(1)}) \cdot \hat{u}^{(1)})\hat{u}^{(1)} + (f(\hat{u}^{(1)}) \cdot \hat{u}^{(2)})\hat{u}^{(2)} \\ &= \left((\sqrt{2}, 0) \cdot \frac{1}{\sqrt{2}}(1, 1)\right)\hat{u}^{(1)} + \left((\sqrt{2}, 0) \cdot \frac{1}{\sqrt{2}}(-1, 1)\right)\hat{u}^{(2)} \\ &= (1)\hat{u}^{(1)} + (-1)\hat{u}^{(2)} = \hat{u}^{(1)} - \hat{u}^{(2)} \end{aligned}$$

Example (Cont.)

- Similarly, for $\vec{f}(\hat{u}^{(2)})$:

$$
\begin{aligned}
f(\hat{u}^{(2)}) &= (f(\hat{u}^{(2)}) \cdot \hat{u}^{(1)})\hat{u}^{(1)} + (f(\hat{u}^{(2)}) \cdot \hat{u}^{(2)})\hat{u}^{(2)} \\
&= \left((0, -\sqrt{2}) \cdot \frac{1}{\sqrt{2}}(1, 1)\right)\hat{u}^{(1)} + \left((0, -\sqrt{2}) \cdot \frac{1}{\sqrt{2}}(-1, 1)\right)\hat{u}^{(2)} \\
&= (-1)\hat{u}^{(1)} + (-1)\hat{u}^{(2)} = -\hat{u}^{(1)} - \hat{u}^{(2)}
\end{aligned}
$$

Solution

- Putting it all together:

$$
\begin{aligned}
f(\vec{x}) &= f(z_1 \hat{u}^{(1)} + z_2 \hat{u}^{(2)}) \\
&= z_1 f(\hat{u}^{(1)}) + z_2 f(\hat{u}^{(2)}) \\
&= z_1 (\hat{u}^{(1)} - \hat{u}^{(2)}) + z_2 (-\hat{u}^{(1)} - \hat{u}^{(2)}) \\
&= (z_1 - z_2)\hat{u}^{(1)} + (-z_1 - z_2)\hat{u}^{(2)}
\end{aligned}
$$

- Or, in coordinate form:

$$
[f(\vec{x})]_{\mathcal{U}} = (z_1 - z_2, -z_1 - z_2)^T
$$

DSC 140B
Representation Learning

Lecture 03 | Part 3

Matrices

Matrices?

- I thought this week was supposed to be about linear algebra… Where are the matrices?

Matrices?

- I thought this week was supposed to be about linear algebra... Where are the matrices?
- What is a matrix, anyways?

What is a matrix?

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix}$$

Recall: Linear Transformations

- A **transformation** $\vec{f}(\vec{x})$ is a function which takes a vector as input and returns a vector of the same dimensionality.

- A transformation $\vec{f}$ is **linear** if

$$\vec{f}(\alpha\vec{u} + \beta\vec{v}) = \alpha\vec{f}(\vec{u}) + \beta\vec{f}(\vec{v})$$

Recall: Linear Transformations

- **Key** consequence of **linearity**: to compute $\vec{f}(\vec{x})$, only need to know what $\vec{f}$ does to basis vectors.

- Example:

$$\vec{x} = 3\hat{e}^{(1)} - 4\hat{e}^{(2)} = \begin{pmatrix} 3 \\ -4 \end{pmatrix}$$

$$\vec{f}(\hat{e}^{(1)}) = -\hat{e}^{(1)} + 3\hat{e}^{(2)}$$

$$\vec{f}(\hat{e}^{(2)}) = 2\hat{e}^{(1)}$$

$$\vec{f}(\vec{x}) =$$

Matrices

- **Idea**: Since $\vec{f}$ is defined by what it does to basis, place $\vec{f}(\hat{e}^{(1)})$, $\vec{f}(\hat{e}^{(2)})$, ... into a table as columns
- This is the **matrix** representing[2] $\vec{f}$

$$\vec{f}(\hat{e}^{(1)}) = -\hat{e}^{(1)} + 3\hat{e}^{(2)} = \begin{pmatrix} -1 \\ 3 \end{pmatrix}$$

$$\vec{f}(\hat{e}^{(2)}) = 2\hat{e}^{(1)} = \begin{pmatrix} 2 \\ 0 \end{pmatrix}$$

$$\begin{pmatrix} -1 & 2 \\ 3 & 0 \end{pmatrix}$$

---

[2] with respect to the standard basis $\hat{e}^{(1)}, \hat{e}^{(2)}$

Example

Write the matrix representing $\vec{f}$ with respect to the standard basis, given:

$$\vec{f}(\hat{e}^{(1)}) = (1, 4, 7)^T$$
$$\vec{f}(\hat{e}^{(2)}) = (2, 5, 8)^T$$
$$\vec{f}(\hat{e}^{(3)}) = (3, 6, 9)^T$$

Exercise

Suppose $\vec{f}$ has the matrix below:

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix}$$

Let $\vec{x} = (-2, 1, 3)^T$. What is $\vec{f}(\vec{x})$?

- A) $(3, 12, 21)^T$
- B) $(-2, 1, 3)^T$
- C) $(6, 15, 24)^T$
- D) $(9, 15, 21)^T$

Main Idea

A square ($n \times n$) matrix can be interpreted as a compact representation of a linear transformation $f : \mathbb{R}^n \to \mathbb{R}^n$.

What is matrix multiplication?

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} \begin{pmatrix} -2 \\ 1 \\ 3 \end{pmatrix} = \begin{pmatrix} \\ \\ \\ \end{pmatrix}$$

A low-level definition

$$(A\vec{x})_i = \sum_{j=1}^{n} A_{ij} x_j$$

A low-level interpretation

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} \begin{pmatrix} -2 \\ 1 \\ 3 \end{pmatrix} = -2 \begin{pmatrix} 1 \\ 4 \\ 7 \end{pmatrix} + 1 \begin{pmatrix} 2 \\ 5 \\ 8 \end{pmatrix} + 3 \begin{pmatrix} 3 \\ 6 \\ 9 \end{pmatrix}$$

In general...

$$\begin{pmatrix} \uparrow & \uparrow & \uparrow \\ \vec{a}^{(1)} & \vec{a}^{(2)} & \vec{a}^{(3)} \\ \downarrow & \downarrow & \downarrow \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = x_1 \vec{a}^{(1)} + x_2 \vec{a}^{(2)} + x_3 \vec{a}^{(3)}$$

Matrix Multiplication

$$\vec{x} = x_1 \hat{e}^{(1)} + x_2 \hat{e}^{(2)} + x_3 \hat{e}^{(3)} = (x_1, x_2, x_3)^T$$

$$\vec{f}(\vec{x}) = x_1 \vec{f}(\hat{e}^{(1)}) + x_2 \vec{f}(\hat{e}^{(2)}) + x_3 \vec{f}(\hat{e}^{(3)})$$

$$A = \begin{pmatrix} \uparrow & \uparrow & \uparrow \\ \vec{f}(\hat{e}^{(1)}) & \vec{f}(\hat{e}^{(2)}) & \vec{f}(\hat{e}^{(3)}) \\ \downarrow & \downarrow & \downarrow \end{pmatrix}$$

$$A\vec{x} = \begin{pmatrix} \uparrow & \uparrow & \uparrow \\ \vec{f}(\hat{e}^{(1)}) & \vec{f}(\hat{e}^{(2)}) & \vec{f}(\hat{e}^{(3)}) \\ \downarrow & \downarrow & \downarrow \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix}$$

$$= x_1 \vec{f}(\hat{e}^{(1)}) + x_2 \vec{f}(\hat{e}^{(2)}) + x_3 \vec{f}(\hat{e}^{(3)})$$

Matrix Multiplication

- Matrix A represents a linear transformation $\vec{f}$
  - With respect to the standard basis
  - If we use a different basis, the matrix changes!

- Matrix multiplication $A\vec{x}$ **evaluates** $\vec{f}(\vec{x})$

What are they, *really*?

- Matrices are sometimes just tables of numbers.
- But they often have a deeper meaning.

Main Idea

A square $(n \times n)$ matrix can be interpreted as a compact representation of a linear transformation $\vec{f} : \mathbb{R}^n \to \mathbb{R}^n$.

What's more, if A represents $\vec{f}$, then $A\vec{x} = \vec{f}(\vec{x})$; that is, multiplying by A is the same as evaluating $\vec{f}$.

Example

$$\vec{x} = 3\hat{e}^{(1)} - 4\hat{e}^{(2)} = \begin{pmatrix} 3 \\ -4 \end{pmatrix}$$

$$\vec{f}(\hat{e}^{(1)}) = -\hat{e}^{(1)} + 3\hat{e}^{(2)}$$

$$\vec{f}(\hat{e}^{(2)}) = 2\hat{e}^{(1)}$$

$$\vec{f}(\vec{x}) =$$

$$A =$$

$$A\vec{x} =$$

Note

- All of this works because we assumed $\vec{f}$ is **linear**.
- If it isn't, evaluating $\vec{f}$ isn't so simple.

Note

- All of this works because we assumed $\vec{f}$ is **linear**.
- If it isn't, evaluating $\vec{f}$ isn't so simple.
- Linear algebra = simple!

Matrices in Other Bases

- The matrix of a linear transformation wrt the **standard basis**:

$$\begin{pmatrix} \uparrow & \uparrow & \uparrow & \\ \vec{f}(\hat{e}^{(1)}) & \vec{f}(\hat{e}^{(2)}) & \cdots & \vec{f}(\hat{e}^{(d)}) \\ \downarrow & \downarrow & \downarrow & \end{pmatrix}$$

- With respect to basis $\mathcal{U}$:

$$\begin{pmatrix} \uparrow & \uparrow & \uparrow & \\ [\vec{f}(\hat{u}^{(1)})]_{\mathcal{U}} & [\vec{f}(\hat{u}^{(2)})]_{\mathcal{U}} & \cdots & [\vec{f}(\hat{u}^{(d)})]_{\mathcal{U}} \\ \downarrow & \downarrow & \downarrow & \end{pmatrix}$$

Example

- Consider the transformation $\vec{f}$ which "mirrors" a vector over the line of 45°.

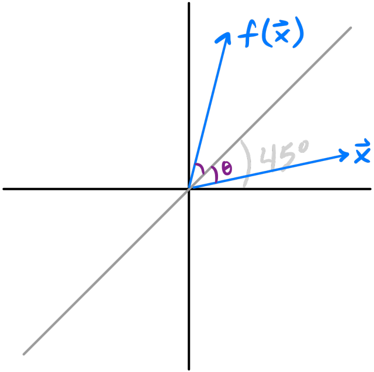

- What is its matrix in the standard basis?

Example

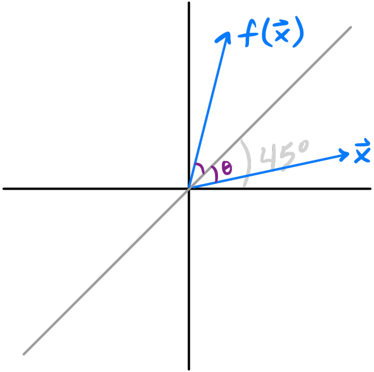

- Let $\hat{u}^{(1)} = \frac{1}{\sqrt{2}}(1, 1)^T$
- Let $\hat{u}^{(2)} = \frac{1}{\sqrt{2}}(-1, 1)^T$
- What is $[\vec{f}(\hat{u}^{(1)})]_{\mathcal{U}}$?
- $[\vec{f}(\hat{u}^{(2)})]_{\mathcal{U}}$?
- What is the matrix?

DSC 140B

Representation Learning

Lecture 03 | Part 4

The Spectral Theorem

Eigenvectors

- Let A be an $n \times n$ matrix. An **eigenvector** of A with **eigenvalue** λ is a nonzero vector $\vec{v}$ such that $A\vec{v} = \lambda\vec{v}$.

Eigenvectors (of Linear Transformations)

- Let $\vec{f}$ be a linear transformation. An **eigenvector** of $\vec{f}$ with **eigenvalue** λ is a nonzero vector $\vec{v}$ such that $f(\vec{v}) = \lambda \vec{v}$.

Importance

- We will see why eigenvectors are important in the next part.
- For now: what are they?

Geometric Interpretation

- Recall: $\vec{v}$ is an eigenvector if $\vec{f}(\vec{v}) = \lambda\vec{v}$.
- Meaning: when $\vec{f}$ is applied to one of its eigenvectors, $\vec{f}$ simply scales it.

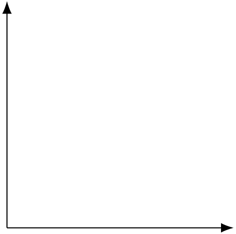

Geometric Interpretation

- The eigenvalue, λ, tells us how much the eigenvector is scaled.
  - If $\lambda > 1$, the eigenvector is stretched.
  - If $0 < \lambda < 1$, the eigenvector is shrunk.
  - If $\lambda < 0$, the eigenvector is flipped and scaled.

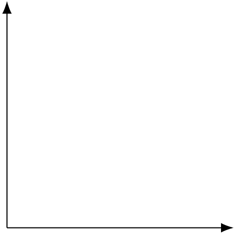

Exercise

Draw as many (linearly independent) eigenvectors as you can:

$$A = \begin{pmatrix} 5 & 0 \\ 0 & 2 \end{pmatrix}$$

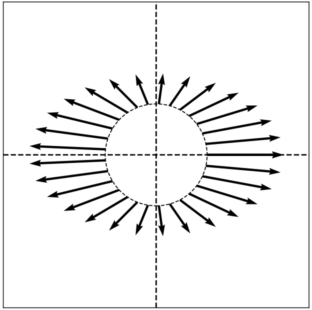

Finding Eigenvectors

- We typically compute the eigenvectors of a matrix with a computer.
- But it can help our understanding to find them "graphically".

Procedure

Given a matrix A (or transformation $\vec{f}$), to find an eigenvector "graphically".

1. Think about (or draw) the output of $\vec{f}$ for a handful of unit vector inputs.
   - Linear transformations are continuous so you can "interpolate".

2. Find place(s) where the input vector and the output vector are parallel.

Exercise

Draw as many (linearly independent) eigenvectors as you can.

$$A = \begin{pmatrix} 2 & -1 \\ -1 & 3 \end{pmatrix}$$

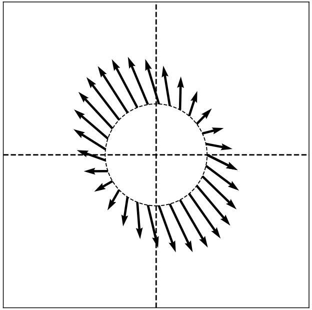

Exercise

Consider the linear transformation which mirrors its input over the line of 45°. Give two orthogonal eigenvectors of the transformation.

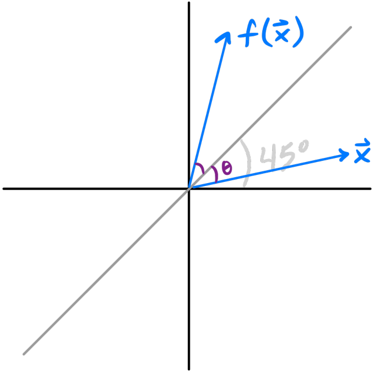

Exercise

Draw as many (linearly independent) eigenvectors as you can.

$$A = \begin{pmatrix} 5 & 5 \\ -10 & 12 \end{pmatrix}$$

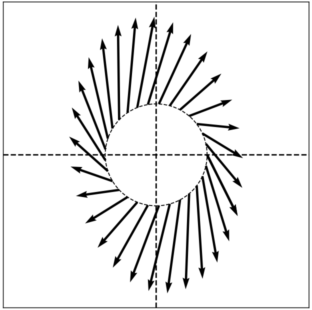

Caution!

- Not all matrices have even one eigenvector![3]
- When does a matrix have multiple (linearly independent) eigenvectors?

[3] That is, with a *real-valued* eigenvalue.

Example Linear Transformation

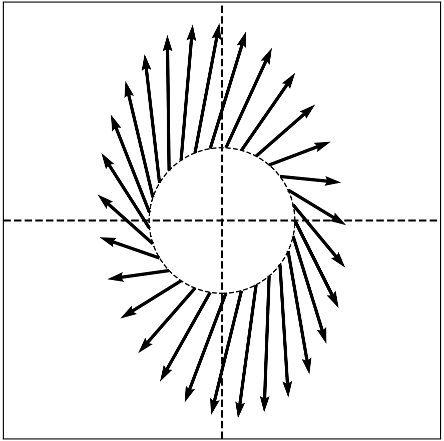

$$A = \begin{pmatrix} 5 & 5 \\ -10 & 12 \end{pmatrix}$$

Example Linear Transformation

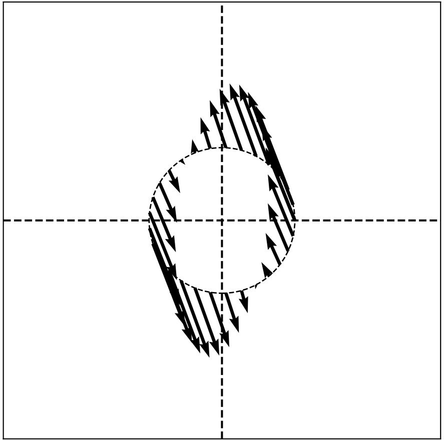

$$A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

Example Linear Transformation

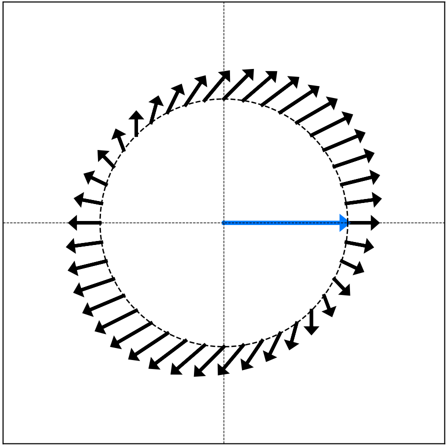

$$A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$$

Example Symmetric Linear Transformation

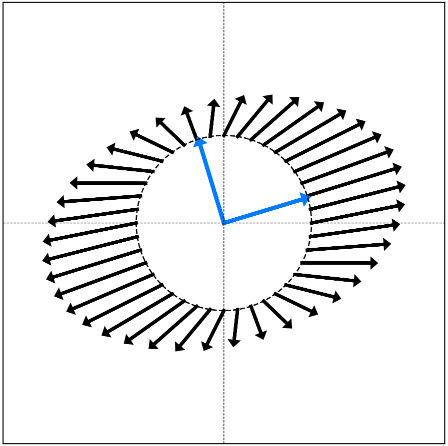

$$A = \begin{pmatrix} 2 & -1 \\ -1 & 3 \end{pmatrix}$$

Example Symmetric Linear Transformation

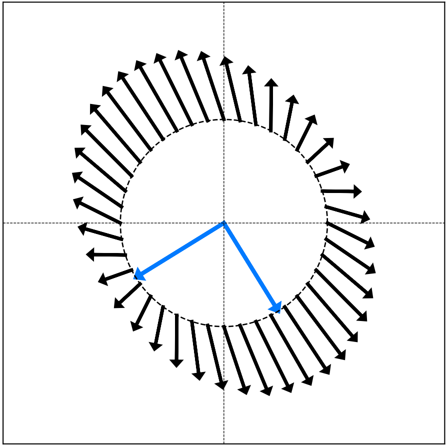

$$A = \begin{pmatrix} 2 & -1 \\ -1 & 3 \end{pmatrix}$$

Example Symmetric Linear Transformation

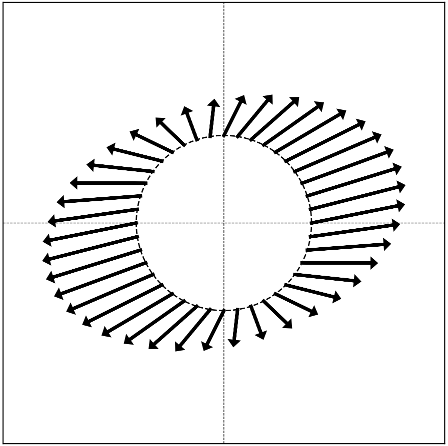

$$A = \begin{pmatrix} 5 & 1 \\ 1 & 2 \end{pmatrix}$$

Example Symmetric Linear Transformation

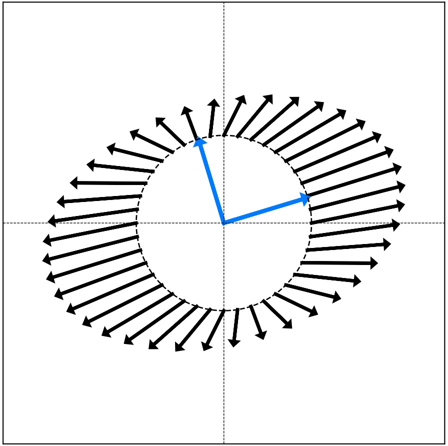

$$A = \begin{pmatrix} 5 & 1 \\ 1 & 2 \end{pmatrix}$$

Observation

- It seems that there is something special about **symmetric** matrices...

Symmetric Matrices

- Recall: a matrix A is **symmetric** if $A^T = A$.

$$\begin{pmatrix} 1 & 2 & 3 \\ 2 & 4 & 5 \\ 3 & 5 & 6 \end{pmatrix}$$

- A linear transformation $\vec{f}$ is **symmetric** if its matrix representation is symmetric.

Observation #1

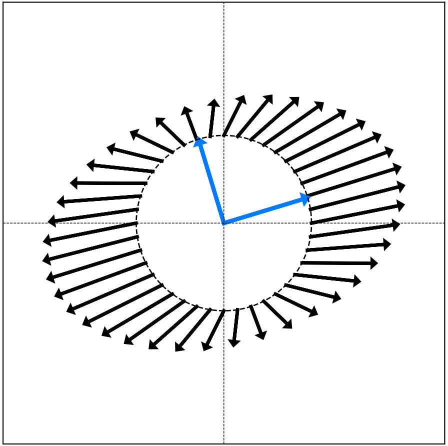

- Symmetric linear transformations have **axes of symmetry**.
  - One for each dimension.

Observation #2

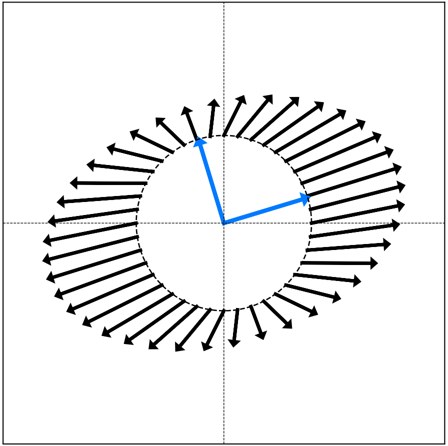

- The axes of symmetry are **orthogonal** to one another.

Observation #3

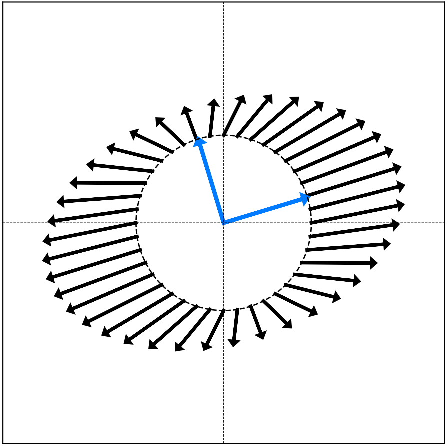

- The action of $\vec{f}$ along an axis of symmetry is simply to **scale** its input.
- That is, the **eigenvectors** point along the axes of symmetry.

Observation #4

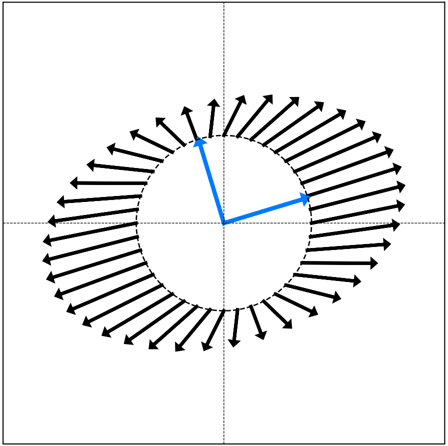

- The size of this scaling can be different for each axis.

Main Idea

The **eigenvectors** of a symmetric linear transformation (matrix) are its axes of symmetry. The **eigenvalues** describe how much each axis of symmetry is scaled.

Exercise

Consider the linear transformation which mirrors its input over the line of 45°. Give two orthogonal eigenvectors of the transformation.

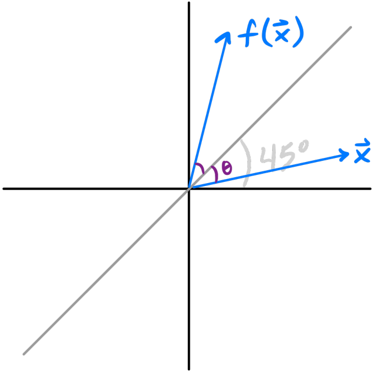

How many?

- The symmetric 2×2 matrices we saw all had 2 orthogonal eigenvectors.
- Does a 3×3 symmetric matrix have 3 orthogonal eigenvectors?
- What about $n \times n$ symmetric matrices?

The Spectral Theorem[4]

Theorem

Let A be an $n \times n$ **symmetric** matrix. Then you can always find n eigenvectors of A which are all mutually orthogonal.

[4] for symmetric matrices

Careful!

- The spectral theorem *does not* say that an $n \times n$ matrix has n eigenvectors!

Exercise

Consider the 2×2 identity matrix. How many (unit) eigenvectors does it have?

A. 0

B. 1

C. 2

D. ∞

Solution

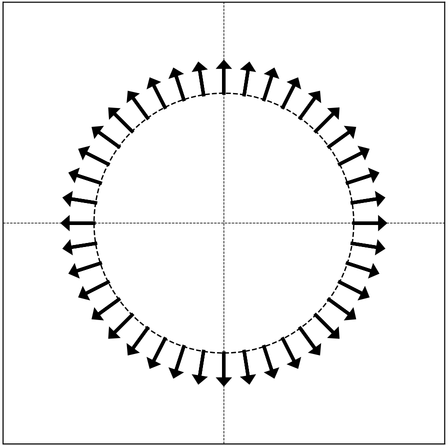

- Infinitely many!
- *Every* (nonzero) vector is an eigenvector with eigenvalue 1.

Solution

- It would be incorrect to say that the identity matrix has just 2 orthogonal eigenvectors.

- Instead, the spectral theorem says: “You can find 2 different orthogonal eigenvectors of I.”

- There are infinitely-many ways to do this!
  - $(1, 0)^T$ and $(0, 1)^T$
  - $(1/\sqrt{2}, 1/\sqrt{2})^T$ and $(-1/\sqrt{2}, 1/\sqrt{2})^T$
  - etc.

Diagonal Matrices

- If A is diagonal, its eigenvectors are simply the standard basis vectors.

$$A = \begin{pmatrix} 2 & 0 \\ 0 & 5 \end{pmatrix}$$

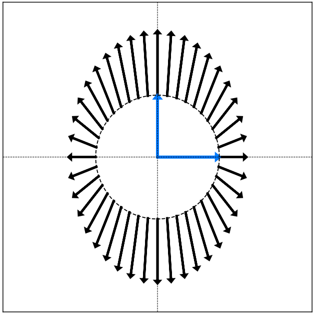

Off-diagonal elements

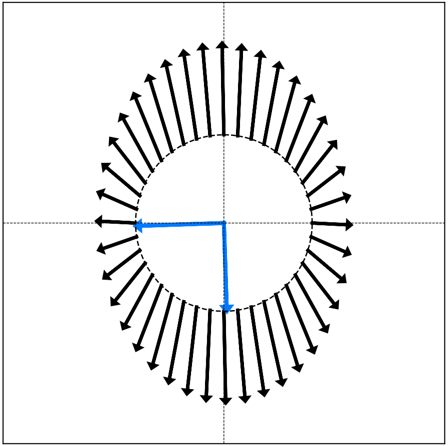

$$A = \begin{pmatrix} 2 & -0.1 \\ -0.1 & 5 \end{pmatrix}$$

Off-diagonal elements

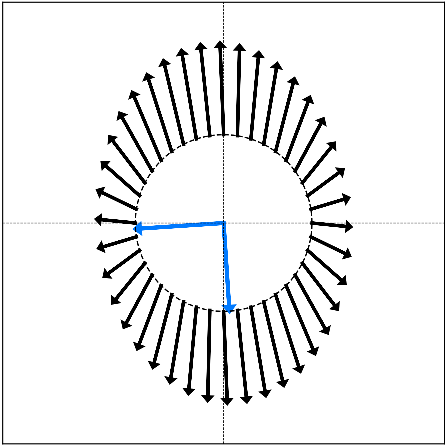

$$A = \begin{pmatrix} 2 & -0.2 \\ -0.2 & 5 \end{pmatrix}$$

Off-diagonal elements

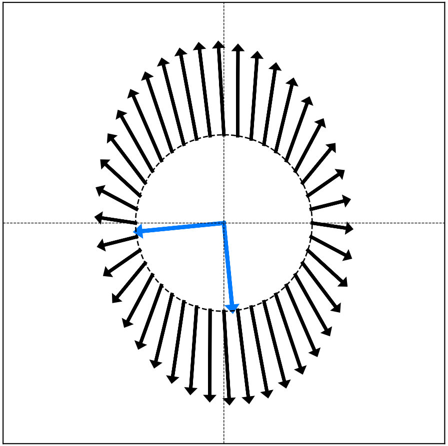

$$A = \begin{pmatrix} 2 & -0.3 \\ -0.3 & 5 \end{pmatrix}$$

Off-diagonal elements

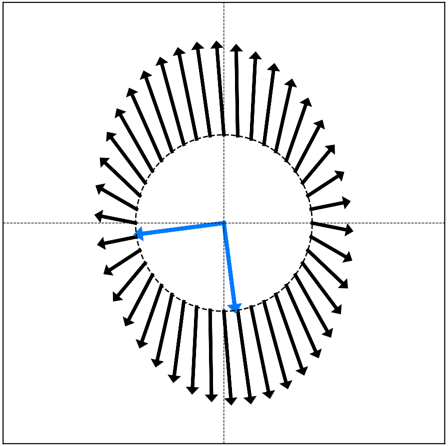

$$A = \begin{pmatrix} 2 & -0.4 \\ -0.4 & 5 \end{pmatrix}$$

Off-diagonal elements

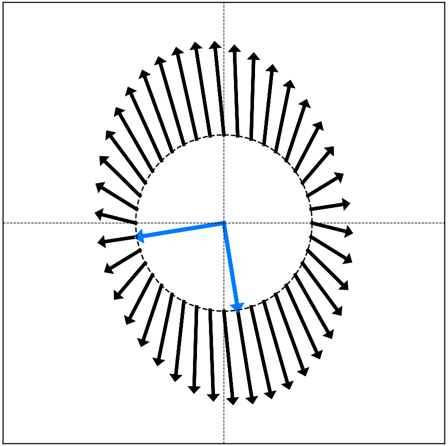

$$A = \begin{pmatrix} 2 & -0.5 \\ -0.5 & 5 \end{pmatrix}$$

Off-diagonal elements

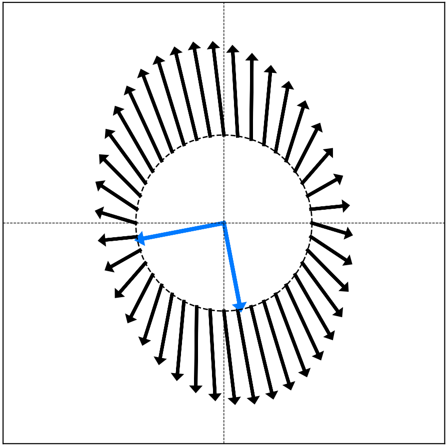

$$A = \begin{pmatrix} 2 & -0.6 \\ -0.6 & 5 \end{pmatrix}$$

Off-diagonal elements

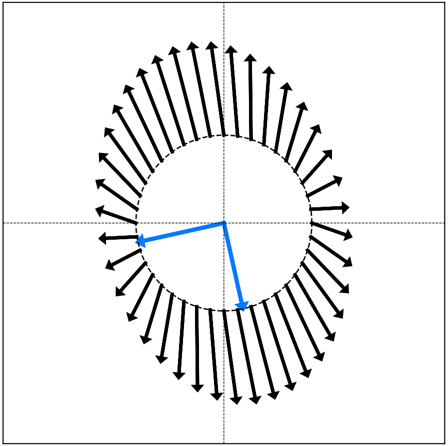

$$A = \begin{pmatrix} 2 & -0.7 \\ -0.7 & 5 \end{pmatrix}$$

Off-diagonal elements

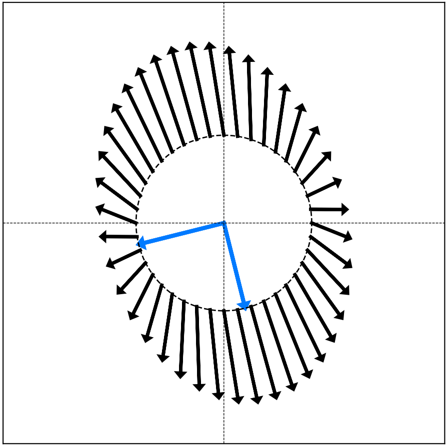

$$A = \begin{pmatrix} 2 & -0.8 \\ -0.8 & 5 \end{pmatrix}$$

Off-diagonal elements

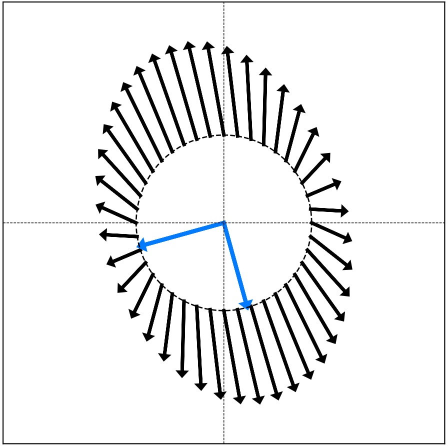

$$A = \begin{pmatrix} 2 & -0.9 \\ -0.9 & 5 \end{pmatrix}$$

Non-Diagonal Symmetric Matrices

- When a symmetric matrix is not diagonal, its eigenvectors are not the standard basis vectors.
- But they are still orthogonal!

Computing Eigenvectors

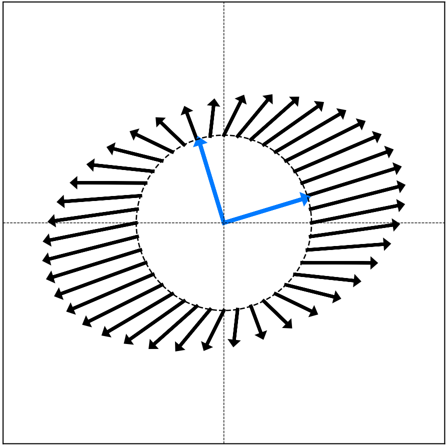

Use `np.linalg.eigh`[a]:

```
>> A = np.array([[2, -1], [-1, 3]])
>> np.linalg.eigh(A)
(array([1.38196601, 3.61803399]),
 array([[-0.85065081, -0.52573111],
        [-0.52573111,  0.85065081]]))
```

---

[a] if the input is *symmetric*