DSC 190

Machine Learning: Representations

Lecture 7 | Part 1

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 $f(\vec{v}) = \lambda\vec{v}$.

Geometric Interpretation

- When $\vec{f}$ is applied to one of its eigenvectors, $\vec{f}$ simply scales it.

- That is, it doesn't **rotate** it.

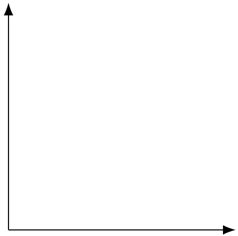

Symmetric Matrices

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

The Spectral Theorem[1]

- **Theorem**: Let A be an $n \times n$ *symmetric* matrix. Then there exist n eigenvectors of A which are all mutually orthogonal.

[1] for symmetric matrices

What?

- What does the spectral theorem mean?
- What is an eigenvector, really?
- Why are they useful?

Example Linear Transformation

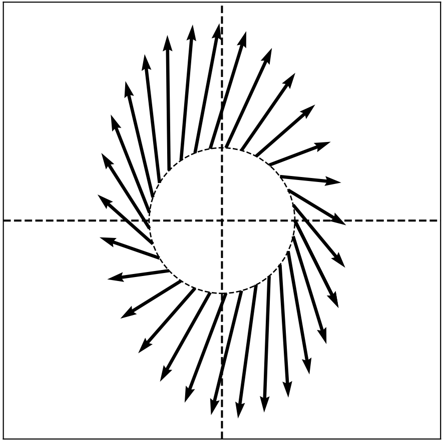

$$A = \begin{pmatrix} 5 & 5 \\ -10 & 12 \end{pmatrix}$$

Example Linear Transformation

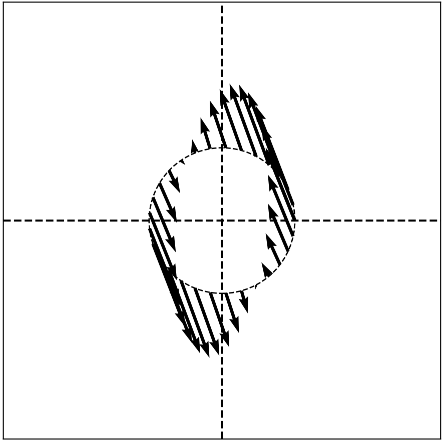

$$A = \begin{pmatrix} -2 & -1 \\ -5 & 3 \end{pmatrix}$$

Example Symmetric Linear Transformation

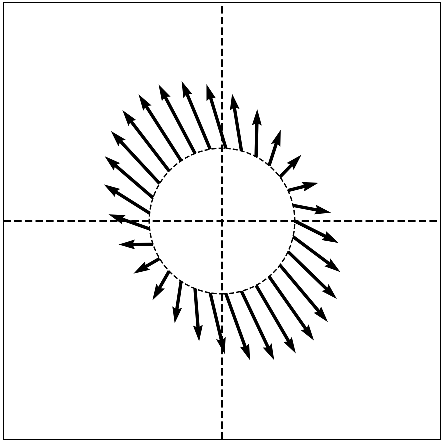

$$A = \begin{pmatrix} 2 & -1 \\ -1 & 3 \end{pmatrix}$$

Example Symmetric Linear Transformation

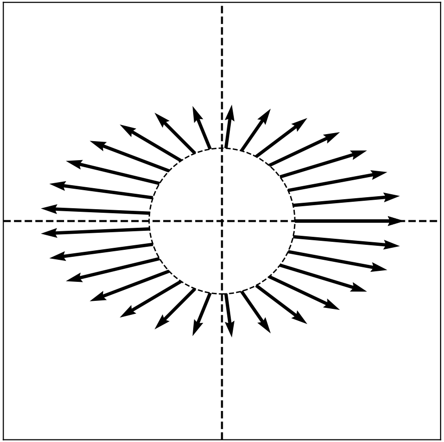

$$A = \begin{pmatrix} 5 & 0 \\ 0 & 2 \end{pmatrix}$$

Observation #1

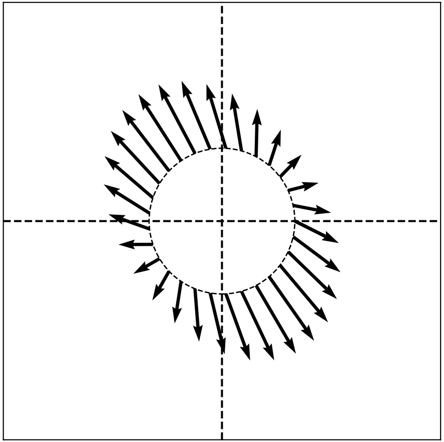

- Symmetric linear transformations have **axes of symmetry**.

Observation #2

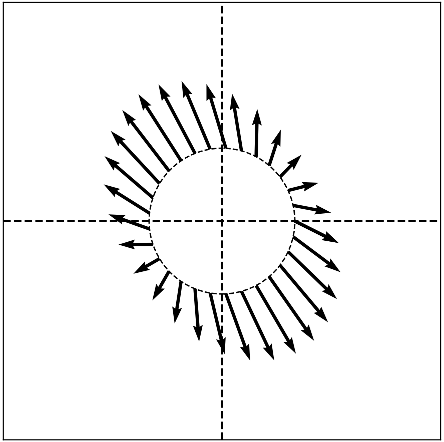

- The axes of symmetry are **orthogonal** to one another.

Observation #3

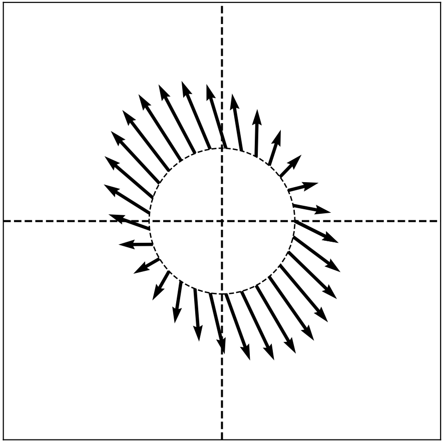

- The action of $\vec{f}$ along an axis of symmetry is simply to scale its input.

Observation #4

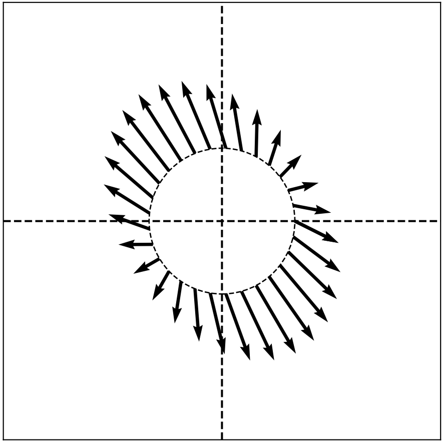

- The size of this scaling can be different for each axis.

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.

Example

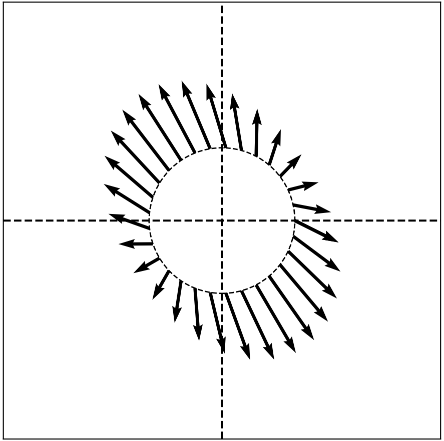

```
>>> A = np.array([[2, -1], [-1, 3]])
>>> np.linalg.eigh(A)
(array([1.38196601, 3.61803399]),
 array([[-0.85065081, -0.52573111],
        [-0.52573111,  0.85065081]]))
```

Off-diagonal elements

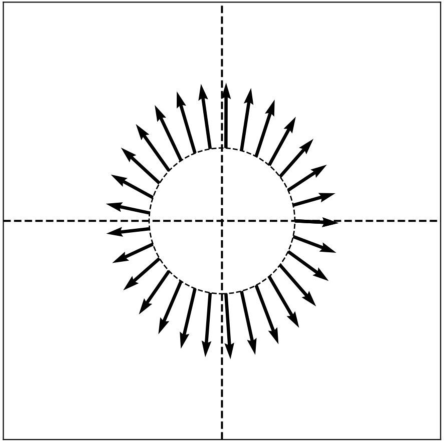

$$A = \begin{pmatrix} 5 & -0.1 \\ -0.1 & 2 \end{pmatrix}$$

Off-diagonal elements

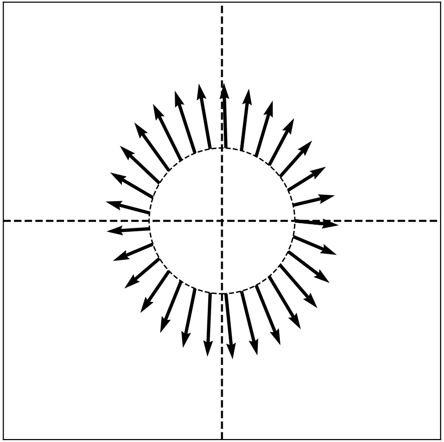

$$A = \begin{pmatrix} 5 & -0.2 \\ -0.2 & 2 \end{pmatrix}$$

Off-diagonal elements

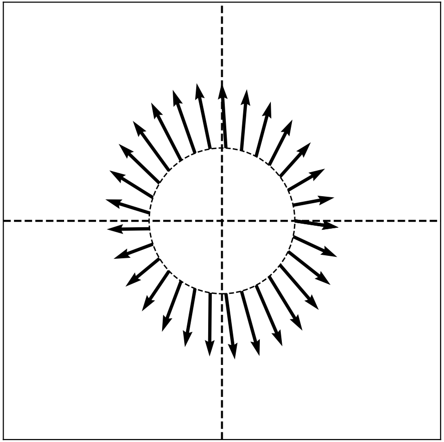

$$A = \begin{pmatrix} 5 & -0.3 \\ -0.3 & 2 \end{pmatrix}$$

Off-diagonal elements

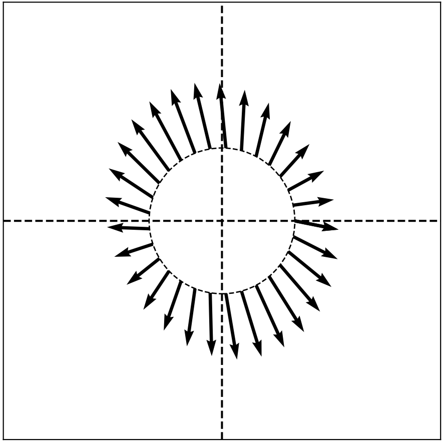

$$A = \begin{pmatrix} 5 & -0.4 \\ -0.4 & 2 \end{pmatrix}$$

Off-diagonal elements

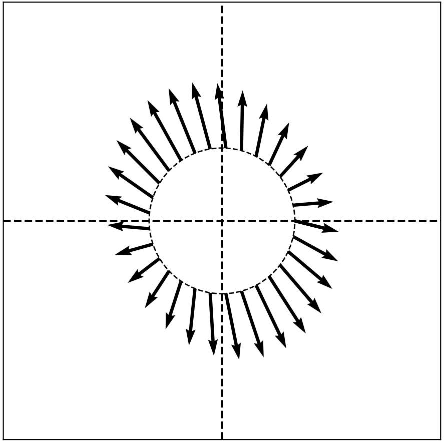

$$A = \begin{pmatrix} 5 & -0.5 \\ -0.5 & 2 \end{pmatrix}$$

Off-diagonal elements

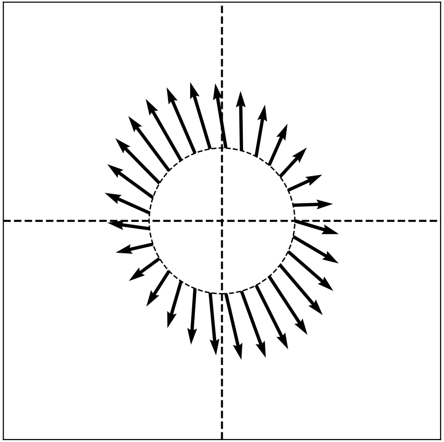

$$A = \begin{pmatrix} 5 & -0.6 \\ -0.6 & 2 \end{pmatrix}$$

Off-diagonal elements

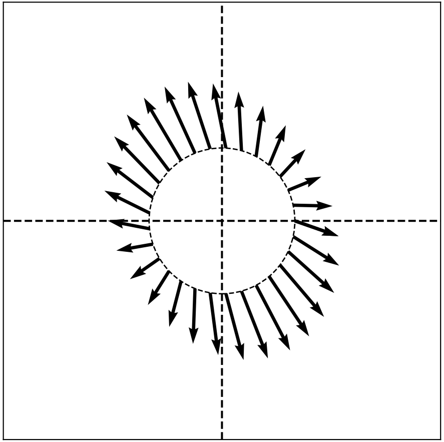

$$A = \begin{pmatrix} 5 & -0.7 \\ -0.7 & 2 \end{pmatrix}$$

Off-diagonal elements

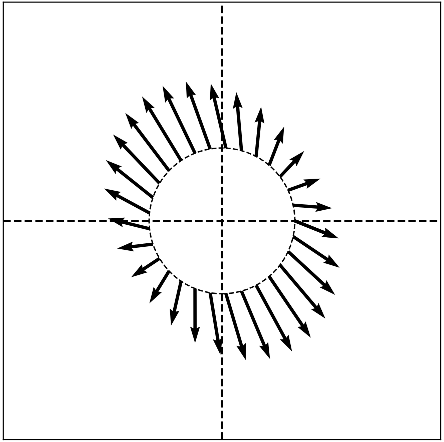

$$A = \begin{pmatrix} 5 & -0.8 \\ -0.8 & 2 \end{pmatrix}$$

Off-diagonal elements

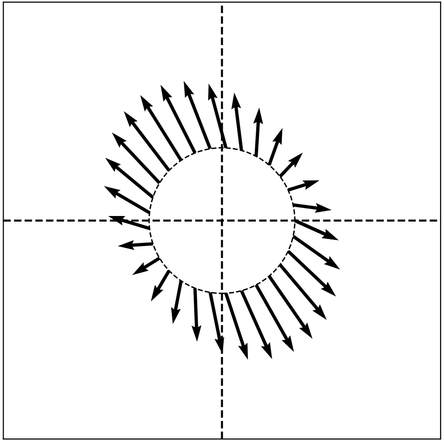

$$A = \begin{pmatrix} 5 & -0.9 \\ -0.9 & 2 \end{pmatrix}$$

Why does $A^T = A$ result in symmetry?

- $A^T = A \implies \vec{f}(\hat{e}^{(1)}) \cdot \hat{e}^{(2)} = \vec{f}(\hat{e}^{(2)}) \cdot \hat{e}^{(1)}$

The Spectral Theorem[2]

- **Theorem**: Let A be an $n \times n$ *symmetric* matrix. Then there exist n eigenvectors of A which are all mutually orthogonal.

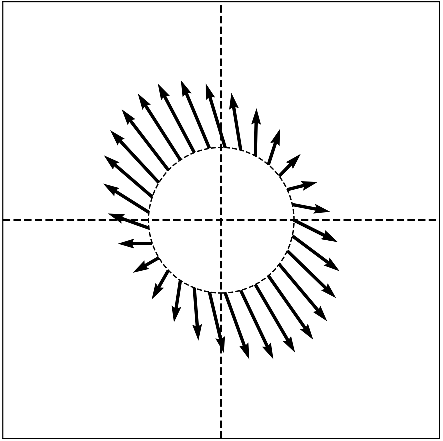

[2]for symmetric matrices

What about total symmetry?

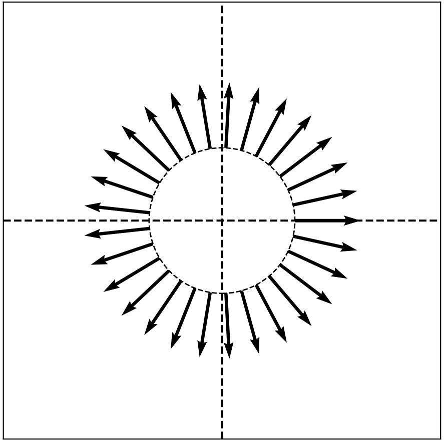

- Every vector is an eigenvector.

$$A = \begin{pmatrix} 3 & 0 \\ 0 & 3 \end{pmatrix}$$

DSC 190

Machine Learning: Representations

Lecture 7 | Part 2

Why are eigenvectors useful?

OK, but why are eigenvectors[3] useful?

- Eigenvectors are nice “building blocks” (basis vectors).
- Eigenvectors are **maximizers** (or minimizers).
- Eigenvectors are **equilibria**.

[3] of symmetric matrices

Eigendecomposition

- Any vector $\vec{x}$ can be written in terms of the eigenvectors of a symmetric matrix.
- This is called its **eigendecomposition**.

Observation #1

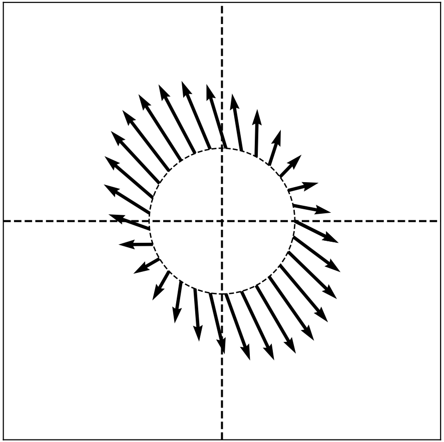

- $\vec{f}(\vec{x})$ is longest along the "main" axis of symmetry.
    - In the direction of the eigenvector with largest eigenvalue.

Main Idea

To maximize $\|\vec{f}(\vec{x})\|$ over unit vectors, pick $\vec{x}$ to be an eigenvector of $\vec{f}$ with the largest eigenvalue (in abs. value).

Main Idea

To minimize $\|\vec{f}(\vec{x})\|$ over unit vectors, pick $\vec{x}$ to be an eigenvector of $\vec{f}$ with the smallest eigenvalue (in abs. value).

Proof

Show that the maximizer of $\|A\vec{x}\|$ s.t., $\|\vec{x}\| = 1$ is the top eigenvector of A.

Corollary

To maximize $\vec{x} \cdot A\vec{x}$ over unit vectors, pick $\vec{x}$ to be top eigenvector of A.

Example

- Maximize $4x_1^2 + 2x_2 + 3x_1x_2$ subject to $x_1^2 + x_2^2 = 1$

Observation #2

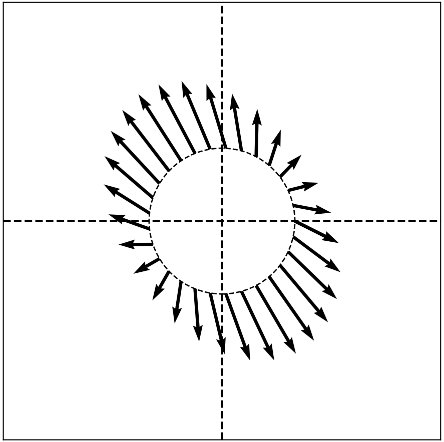

- $\vec{f}(\vec{x})$ rotates $\vec{x}$ towards the “top” eigenvector $\vec{v}$.
- $\vec{v}$ is an equilibrium.

The Power Method

- Method for computing the top eigenvector/value of A.
- Initialize $\vec{x}^{(0)}$ randomly
- Repeat until convergence:
  - Set $\vec{x}^{(i+1)} = A\vec{x}^{(i)} / \|A\vec{x}^{(i)}\|$

DSC 190

Machine Learning: Representations

Lecture 7 | Part 3

Diagonalization

Spectral Theorem (Again)

- **Theorem**: Let A be an $n \times n$ *symmetric* matrix. Then there exists an orthogonal matrix U and a diagonal matrix Λ such that $A = U^T \Lambda U$.

- The *rows* of U are the eigenvectors of A, and the entries of Λ are its eigenvalues.

- U is said to **diagonalize** A.

Note about Bases

- To write the matrix representation of f, you must first choose a basis.

- If it isn't stated, we'll assume the standard basis.

- But we can also write a matrix representing f in some other basis.

$$f(\hat{u}^{(1)}) = 2\hat{u}^{(1)} + 3\hat{u}^{(2)} = (2, 3)^T_{\mathcal{U}}$$
$$f(\hat{u}^{(2)}) = -5\hat{u}^{(1)} - \hat{u}^{(2)} = (-5, -1)^T_{\mathcal{U}}$$

$$A_{\mathcal{U}} =$$

Eigenbasis

- A basis of eigenvectors is particularly natural.
- Example: $\vec{f}(\vec{v}^{(1)}) = \lambda_1 \vec{v}^{(1)}$, $\vec{f}(\vec{v}^{(2)}) = \lambda_2 \vec{v}^{(2)}$
- Matrix representing $\vec{f}$ in the eigenbasis:

Two Approaches

- Approach 1:
  - Write matrix for A w.r.t. standard basis
  - $\vec{f}(\vec{x}) = A\vec{x}$

- Approach 2:
  - Change basis to **eigenbasis**
  - Apply matrix representing $\vec{f}$ in the eigenbasis (simple)
  - Change basis back to original basis

Spectral Theorem (Again)

- **Theorem**: Let A be an $n \times n$ *symmetric* matrix. Then there exists an orthogonal matrix U and a diagonal matrix Λ such that $A = U^T \Lambda U$.

- Interpretation:
  - Change basis by multiplying by U
  - Λ is the representation of $\vec{f}$ in the eigenbasis
  - Change basis back by multiplying by U^T

Geometric Interpretation of $\vec{u} \cdot \vec{v}$

- $\vec{u} \cdot \vec{v} = \|\vec{u}\| \|\vec{v}\| \cos\theta.$

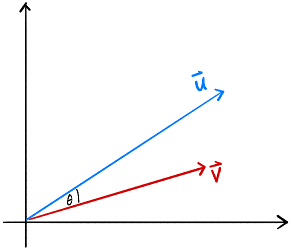

Change of Basis

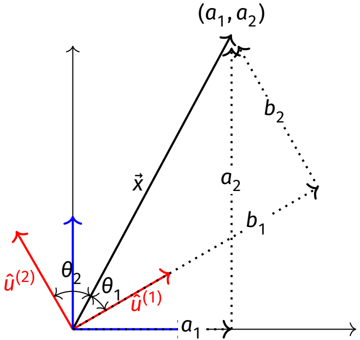

$$\vec{x} = a_1 \hat{e}^{(1)} + a_2 \hat{e}^{(2)}$$
$$\vec{x} = b_1 \hat{u}^{(1)} + b_2 \hat{u}^{(2)}$$

Change of Basis

- Suppose $\hat{u}^{(1)}$ and $\hat{u}^{(2)}$ are our new, **orthonormal** basis vectors.
- We know $\vec{x} = x_1 \hat{e}^{(1)} + x_2 \hat{e}^{(2)}$
- We want to write $\vec{x} = b_1 \hat{u}^{(1)} + b_2 \hat{u}^{(2)}$
- Solution

$$b_1 = \vec{x} \cdot \hat{u}^{(1)} \qquad b_2 = \vec{x} \cdot \hat{u}^{(2)}$$

Example

$$\hat{u}^{(1)} = (\sqrt{3}/2, 1/2)^T$$
$$\hat{u}^{(2)} = (-1/2, \sqrt{3}/2)^T$$
$$\vec{x} = (1/2, 1)^T$$

Change of Basis Matrix

- Changing basis is a linear transformation

$$f(\vec{x}) = (\vec{x} \cdot \hat{u}^{(1)})\hat{u}^{(1)} + (\vec{x} \cdot \hat{u}^{(2)})\hat{u}^{(2)} = \begin{pmatrix} \vec{x} \cdot \hat{u}^{(1)} \\ \vec{x} \cdot \hat{u}^{(2)} \end{pmatrix}_{\mathcal{U}}$$

- We can represent it with a matrix

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

Example

$$\begin{aligned}
\hat{u}^{(1)} &= (\sqrt{3}/2, 1/2)^T \\
\hat{u}^{(2)} &= (-1/2, \sqrt{3}/2)^T \\
f(\hat{e}^{(1)}) &= \\
f(\hat{e}^{(2)}) &= \\
A &=
\end{aligned}$$

Change of Basis Matrix

- Multiplying by this matrix gives the coordinate vector w.r.t. the new basis.
- Example:

$$
\begin{aligned}
\hat{u}^{(1)} &= (\sqrt{3}/2, 1/2)^T \\
\hat{u}^{(2)} &= (-1/2, \sqrt{3}/2)^T \\
A &= \begin{pmatrix} \sqrt{3}/2 & 1/2 \\ -1/2 & \sqrt{3}/2 \end{pmatrix} \\
\vec{x} &= (1/2, 1)^T
\end{aligned}
$$

Change to Eigenbasis

- It can be shown that the matrix which changes basis to the eigenbasis of A is the orthogonal matrix U, whose rows are the eigenvectors of A.