Homework 10: Eigenvalues and Eigenvectors

EECS 245, Fall 2025 at the University of Michigan

due Monday, November 24th, 2025 at 11:59PM Ann Arbor Time

Write your solutions to the following problems by either typing them up or handwriting them on another piece of paper. Homeworks are due to Gradescope by 11:59PM on the due date. See the **syllabus** for details on the slip day policy.

Homework will be evaluated not only on the correctness of your answers, but on your ability to present your ideas clearly and logically. You should always explain and justify your conclusions, using sound reasoning. Your goal should be to convince the reader of your assertions. If a question does not require explanation, it will be explicitly stated.

Before proceeding, make sure you're familiar with the **collaboration policy**.

Total Points: 10 + 12 + 10 + 24 + 16 + 16 + 15 + 12 = 115

Advice:

- The problems in this homework are roughly sorted in order of when we introduced the concepts in class. Problems 2 and 3 (and in some ways, Problem 8) only involve Chapter 5.1. Problems 4 onwards will require Chapter 5.2 and Tuesday 11/18's lecture, though you can work on *some* parts of them without that knowledge.
- Repeatedly, you'll be asked to find eigenvalues and eigenvectors. As usual, you're expected to show all of your work. But, you're encouraged to verify your answers by using `np.linalg.eig` in Python. (Resist the urge to use ChatGPT...)

Problem 1: Homework 9 Solutions Review (10 pts)

Review **the solutions to Homework 9**. Pick **two problem parts** (for example, Problem 4c and Problem 7c) from Homework 9 in which your solutions have the most room for improvement, i.e. where they have unsound reasoning, could be significantly more efficient or clearer, etc. Include a screenshot of your solution to each problem part, and in a few sentences, explain what was deficient and how it could be fixed.

Alternatively, if you think one of your solutions is significantly better than the posted one, copy it here and explain why you think it is better. If you didn't do Homework 9, choose two problem parts from it that look challenging to you, and in a few sentences, explain the key ideas behind their solutions in your own words.

Problem 2: Fundamentals (12 pts)

For each 2×2 matrix A below:

(i) Find the characteristic polynomial of A, and use it to find the eigenvalues of A.

(ii) Find one eigenvector for each eigenvalue of A. Verify that each eigenvector is indeed an eigenvector of A by multiplying it by A.

(iii) **By hand (not using Python or Desmos)**, draw a picture (like the one in Chapter 5.1 titled **Visualizing the eigenvectors of A**) with vectors $\vec{v}_1, A\vec{v}_1, \vec{v}_2, A\vec{v}_2$ as arrows (where $\vec{v}_1$ and $\vec{v}_2$ are the eigenvectors you found above).

a) (4 pts) $A = \begin{bmatrix} 3 & 4 \\ 4 & 3 \end{bmatrix}$

b) (4 pts) $A = \begin{bmatrix} 2 & 1 \\ -2 & 8 \end{bmatrix}$

c) (4 pts) $A = \begin{bmatrix} 3 & 0 \\ 0 & 4 \end{bmatrix}$

Problem 3: Rank One Projection Matrices (10 pts)

Consider the unit vector $\vec{u} = \begin{bmatrix} 1/6 \\ 1/6 \\ 3/6 \\ 5/6 \end{bmatrix}$, and the corresponding rank one projection matrix $P = \vec{u}\vec{u}^T$.

a) (3 pts) Show that $\vec{u}$ is an eigenvector of P. What is its corresponding eigenvalue?

b) (3 pts) Show that if $\vec{v}$ is orthogonal to $\vec{u}$, then $\vec{v}$ is an eigenvector of P. What is its corresponding eigenvalue?

c) (4 pts) Find three different **linearly independent** eigenvectors of P, all corresponding to the eigenvalue 0.

(In the terminology of Problem 4 and Chapter 5.2, these eigenvectors form a basis of the eigenspace of P corresponding to eigenvalue 0.)

Problem 4: Algebraic and Geometric Multiplicities (24 pts)

For each matrix below:

(i) Find its characteristic polynomial in factored form.

(ii) State all eigenvalues along with their algebraic multiplicities.

(iii) For each eigenvalue, find a basis for the eigenspace corresponding to that eigenvalue, and state its geometric multiplicity.

Hint: Recall the trace and determinant tricks from Chapter 5.1, and the fact that the determinant of an upper triangular matrix is the product of the diagonal entries. Only one of the six characteristic polynomials requires a lot of work to factor; the others are relatively straightforward.

a) (4 pts) $A = \begin{bmatrix} 3 & 0 & 0 \\ 0 & 4 & 0 \\ 0 & 0 & 4 \end{bmatrix}$

b) (4 pts) $A = \begin{bmatrix} 3 & 1 \\ 0 & 3 \end{bmatrix}$

c) (4 pts) $A = \begin{bmatrix} 5 & 1 & -1 \\ -1 & 5 & 1 \\ 0 & 2 & 4 \end{bmatrix}$

d) (4 pts) $A = \begin{bmatrix} 2 & 0 & 1 \\ 0 & 2 & 1 \\ 0 & 0 & 3 \end{bmatrix}$

e) (4 pts) $A = \begin{bmatrix} 5 & 0 & 0 & 0 \\ 0 & 3 & 1 & 0 \\ 0 & 0 & 3 & 0 \\ 0 & 0 & 0 & 5 \end{bmatrix}$

f) (4 pts) $A = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$

Problem 5: Diagonalization (16 pts)

Recall from Chapter 5.2 that an $n \times n$ matrix A is **diagonalizable** if it can be written as $A = PDP^{-1}$, where P and D are both $n \times n$, D is a diagonal matrix, and P is invertible. Diagonal matrices are particularly easy to work with: if $A = PDP^{-1}$, then

$$A^k = (PDP^{-1})^k = PD\underbrace{P^{-1}P}_{I}D\underbrace{P^{-1}P}_{I}D\underbrace{P^{-1}P}_{I}\cdots PD\underbrace{P^{-1}P}_{I}DP^{-1} = PD^kP^{-1}$$

a) (4 pts) In each statement, fill in the blanks and provide a brief justification. Each answer is more than just one word or number.

(i) A is diagonalizable if and only if it has ____ eigenvectors.

(ii) A is diagonalizable if and only if the geometric multiplicity of each eigenvalue is ____.

b) (12 pts) For each matrix A in **Problem 4**:

- **if** it is diagonalizable, find matrices V and Λ such that $A = V\Lambda V^{-1}$. (As we saw in Chapter 5.2, this matrix is constructed by placing the eigenvectors of A as the columns of V and the eigenvalues of A as the diagonal entries of Λ. You should have already done most of the work for this; this problem is just a matter of organizing your work.)
- **if not**, explain why it is not diagonalizable.

Problem 6: Adjacency Matrices (16 pts)

Consider the matrix $A = \begin{bmatrix} 0.6 & 0.2 & 0.4 \\ 0.3 & 0.7 & 0.2 \\ 0.1 & 0.1 & 0.4 \end{bmatrix}$. A represents the adjacency matrix of a system with three states. (These systems are called **Markov chains**.)

a) (3 pts) Draw the corresponding state diagram for A, using the diagram at the end of Chapter 5.1 as an example. (Label the states 1, 2, and 3.)

b) (4 pts) Diagonalize A by finding matrices V and Λ such that $A = V\Lambda V^{-1}$. Do this by hand, but then include a screenshot of `numpy` code that verifies that you found the correct V and Λ.

c) (3 pts) Compute A^{10} using the diagonalization you found in part **b)**. *Hint: You should **not** have to multiply ten matrices by hand: only three. State what the three matrices are, and then you can use `numpy` to actually multiply them. Include a screenshot of any code you write and its output.*

d) (3 pts) Let $\vec{x}_0 = \begin{bmatrix} 0.6 \\ 0.3 \\ 0.1 \end{bmatrix}$ be an initial state vector. Using `numpy`, compute $A^{10}\vec{x}_0$. Include a screenshot of any code you write and its output.

e) (3 pts) As $k \to \infty$, what does $A^k\vec{x}_0$ converge to, and why? Make sure your answer references the diagonalization you found in part **b)**.

Problem 7: Regularization Returns (15 pts)

In Homework 9, you found that the **regularized** objective function

$$R_{\text{ridge}}(\vec{w}) = \|\vec{y} - X\vec{w}\|^2 + \lambda\|\vec{w}\|^2$$

was minimized at

$$\vec{w}^*_{\text{ridge}} = (X^TX + \lambda I)^{-1}X^T\vec{y}$$

There, we claimed that the matrix $X^TX + \lambda I$ is invertible for all $\lambda \geq 0$, **even if X is not full rank**, but didn't prove it.

We now have the tools to approach the proof.

a) (4 pts) Recall from Chapter 5.2 that a symmetric $n \times n$ matrix A is said to be **positive semidefinite** if $\vec{v}^TA\vec{v} \geq 0$ for all $\vec{v} \in \mathbb{R}^n$. A is positive semidefinite if and only if all of its eigenvalues are non-negative.

Prove that X^TX is positive semidefinite, where X is an $n \times d$ matrix. *Hint: Suppose $\vec{v}_i$ is an eigenvector of X^TX with eigenvalue λ_i. From there, if you get stuck, take a look at* ***this seemingly unrelated proof from Chapter 2.8*** *for inspiration.*

b) (3 pts) Suppose $\vec{v}_i$ is an eigenvector of X^TX with eigenvalue λ_i. Show that $\vec{v}_i$ is also an eigenvector of $X^TX + \lambda I$. What is its corresponding eigenvalue?

c) (3 pts) Putting parts **a)** and **b)** together, why is it guaranteed that $X^TX + \lambda I$ is invertible for all $\lambda > 0$, even if X is not full rank? ($X^TX + \lambda I$ is said to be **positive definite** for all $\lambda > 0$.)

d) (5 pts) Now, let's explore how adding the regularization term $\lambda\|\vec{w}\|^2$ to the objective function affects the shape of the loss surface.
Open the **the supplemental Jupyter Notebook** we've created for Homework 10, which can either be found **here** on DataHub, or **here** in the course GitHub repository.
This problem is **not** autograded. Instead,

(i) Read through the entire walkthrough, all the way through the end of Problem 7d) (not including Problem 7e)).

(ii) In this PDF, include a screenshot of the diagram with a slider, showing that you've moved it all the way to the right, at $\lambda = 100000$.

(iii) In this PDF, include answers to both of the following questions:

- Why is it called ridge regression?
- How does the inclusion of the $\lambda\|\vec{w}\|^2$ term change the **convexity** of the loss surface?

If you'd like to read more about regularization, and **how we actually choose the value of λ in practice**, read more from **EECS 398 here**.

Problem 8: PageRank (12 pts)

This problem involves writing code and submitting it to the Gradescope autograder. The goal of this problem is to allow you to implement Google's PageRank algorithm in code and think through some of its pitfalls and variants.

There are two ways to access the supplemental Jupyter Notebook:

- **Option 1**: Click **here** to open `hw10.ipynb` on DataHub. Before doing so, read the instructions on the **Tech Support** page on how to use the DataHub.
- **Option 2**: Set up a Jupyter Notebook environment locally, use `git` to clone our course repository, and open `homeworks/hw10/hw10.ipynb`. For instructions on how to do this, see the **Tech Support** page of the course website.

This problem is entirely autograded; to receive credit for Problem 8 of this homework, you'll need to submit your completed notebook to the autograder on Gradescope. Your submission time for Homework 10 is the **latter** of your PDF and code submission times.