

Due Monday, 22 Jan 2024, by 11:59pm to Gradescope.

Covers material up to Introduction to machine learning refresher 1.

100 points total.

1. (25 points) **Linear algebra refresher.**

(a) (12 points) Let \mathbf{Q} be a real orthogonal matrix.

i. (3 points) Show that \mathbf{Q}^T and \mathbf{Q}^{-1} are also orthogonal.

ii. (3 points) Show that \mathbf{Q} has eigenvalues with norm 1.

iii. (3 points) Show that the determinant of \mathbf{Q} is either +1 or -1.

iv. (3 points) Show that \mathbf{Q} defines a length preserving transformation.

(b) (8 points) Let \mathbf{A} be a matrix.

i. (4 points) What is the relationship between the singular vectors of \mathbf{A} and the eigenvectors of $\mathbf{A}\mathbf{A}^T$? What about $\mathbf{A}^T\mathbf{A}$?

ii. (4 points) What is the relationship between the singular values of \mathbf{A} and the eigenvalues of $\mathbf{A}\mathbf{A}^T$? What about $\mathbf{A}^T\mathbf{A}$?

(c) (5 points) True or False. Partial credit on an incorrect solution may be awarded if you justify your answer.

i. Every linear operator in an n -dimensional vector space has n distinct eigenvalues.

ii. A non-zero sum of two eigenvectors of a matrix \mathbf{A} is an eigenvector.

iii. If a matrix \mathbf{A} has the positive semidefinite property, i.e., $\mathbf{x}^T\mathbf{A}\mathbf{x} \geq 0$ for all \mathbf{x} , then its eigenvalues must be non-negative.

iv. The rank of a matrix can exceed the number of distinct non-zero eigenvalues.

v. A non-zero sum of two eigenvectors of a matrix \mathbf{A} corresponding to the same eigenvalue λ is always an eigenvector.

2. (25 points) **Probability refresher.**

(a) (5 points) A and B are involved in a duel. The rules of the duel are that they are to pick up their guns and shoot at each other simultaneously. If one or both are hit, then the duel is over. If both shots miss, then they repeat the process. Suppose that the results of the shots are independent and that each shot of A will hit B with probability p_A and each shot of B will hit A with probability p_B . What is:

i. (1 point) the probability that A is not hit?

ii. (1 point) the probability that both duelists are hit?

iii. (1 point) the probability that the duel ends after the n^{th} round of shots?

- iv. (1 point) the conditional probability that the duel ends after the n^{th} round of shots given that A is not hit?
- v. (1 point) the conditional probability that the duel ends after the n^{th} round of shots given that both duelists are hit?
- (b) (5 points) Consider a group of 18 UCLA faculty members containing 6 ECE faculty, 6 CSE faculty, and 6 Math faculty. We mix up the faculty members and place all 18 faculties around a circular table, with one faculty member per seat.
- (2 points) A faculty member is called “isolated” if his/her department does not agree with either of the nearby faculties (i.e., if he/she has a different department than the faculty to his/her right and a different department than the faculty to his/her left). Let X denote the number of isolated faculties. Find $E(X)$.
 - (2 points) A faculty member is called “semi-happy” if his/her department agrees with exactly one (but not both) of the nearby faculties (i.e., if his/her department agrees with the department of the faculty on his/her left or on his/her right, but not both). Let Y denote the number of semi-happy faculties. Find $E(Y)$.
 - (1 point) A faculty member is called “joyous” if his/her department agrees with both of the nearby faculties (i.e., if his/her department agrees with the department of the faculties on his/her left and on his/her right). Let Z denote the number of joyous faculties. Find $E(Z)$.
- (c) (5 points) There is a screening test for lung cancer that looks at the level of LSA (lung specific antigen) in the blood. There are a number of reasons besides lung cancer that a man can have elevated LSA levels. In addition, many types of lung cancer develop so slowly that they are never a problem. Unfortunately, there is currently no test to distinguish the different types and using the test is controversial because it’s hard to quantify the accuracy rates and the harm done by false positives. For this problem, we will call a positive test a true positive if it catches a dangerous type of lung cancer. Also, we will assume the following numbers:
- Rate of dangerous type of lung cancer among men over 30 = 0.0005
 - True positive rate for the test = 0.9
 - False positive rate for the test = 0.01
- Suppose you randomly select a man over 30 and perform a screening test.
- (3 points) What is the probability that the man has a dangerous type of the disease given that he had a positive test?
 - (2 points) What is the probability that the man has a dangerous type of the disease given that he had a negative test?
- (d) (5 points) Let x_1, x_2, \dots, x_n be identically distributed random variables. A random vector, \mathbf{x} , is defined as

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}$$

What is $\mathbb{E}(\mathbf{Ax} + \mathbf{b})$ in terms of $\mathbb{E}(\mathbf{x})$, given that \mathbf{A} and \mathbf{b} are deterministic?

(e) (5 points) Let

$$\mathbf{cov}(\mathbf{x}) = \mathbb{E}((\mathbf{x} - \mathbb{E}\mathbf{x})(\mathbf{x} - \mathbb{E}\mathbf{x})^T)$$

What is $\mathbf{cov}(\mathbf{Ax} + \mathbf{b})$ in terms of $\mathbf{cov}(\mathbf{x})$, given that \mathbf{A} and \mathbf{b} are deterministic?

3. (10 points) **Multivariate derivatives.**

(a) (1 points) Let $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{y} \in \mathbb{R}^m$, and $\mathbf{A} \in \mathbb{R}^{n \times m}$. What is $\nabla_{\mathbf{x}} \mathbf{x}^T \mathbf{A} \mathbf{y}$?

(b) (1 points) Let $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{y} \in \mathbb{R}^m$, and $\mathbf{A} \in \mathbb{R}^{n \times m}$. What is $\nabla_{\mathbf{y}} \mathbf{x}^T \mathbf{A} \mathbf{y}$?

(c) (1 points) Let $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{y} \in \mathbb{R}^m$, and $\mathbf{A} \in \mathbb{R}^{n \times m}$. What is $\nabla_{\mathbf{A}} \mathbf{x}^T \mathbf{A} \mathbf{y}$?

(d) (1 points) Let $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{A} \in \mathbb{R}^{n \times n}$, and let $f = \mathbf{x}^T \mathbf{A} \mathbf{x} + \mathbf{b}^T \mathbf{x}$. What is $\nabla_{\mathbf{x}} f$?

(e) (1 points) Let $\mathbf{A} \in \mathbb{R}^{n \times n}$, $\mathbf{B} \in \mathbb{R}^{n \times n}$ and $f = \text{tr}(\mathbf{AB})$. What is $\nabla_{\mathbf{A}} f$?

(f) (2 points) Let $\mathbf{A} \in \mathbb{R}^{n \times n}$, $\mathbf{B} \in \mathbb{R}^{n \times n}$ and $f = \text{tr}(\mathbf{BA} + \mathbf{A}^T \mathbf{B} + \mathbf{A}^2 \mathbf{B})$. What is $\nabla_{\mathbf{A}} f$?

(g) (3 points) Let $\mathbf{A} \in \mathbb{R}^{n \times n}$, $\mathbf{B} \in \mathbb{R}^{n \times n}$ and $f = \|\mathbf{A} + \lambda \mathbf{B}\|_F^2$. What is $\nabla_{\mathbf{A}} f$?

4. (10 points) **Deriving least-squares with matrix derivatives.**

In least-squares, we seek to estimate some multivariate output \mathbf{y} via the model

$$\hat{\mathbf{y}} = \mathbf{W} \mathbf{x}$$

In the training set we're given paired data examples $(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})$ from $i = 1, \dots, n$. Least-squares is the following quadratic optimization problem:

$$\min_{\mathbf{W}} \frac{1}{2} \sum_{i=1}^n \|\mathbf{y}^{(i)} - \mathbf{W} \mathbf{x}^{(i)}\|^2$$

Derive the optimal \mathbf{W} .

Where \mathbf{W} is a matrix, and for each example in the training set, both $\mathbf{x}^{(i)}$ and $\mathbf{y}^{(i)} \forall i = 1, \dots, n$ are vectors.

Hint: you may find the following derivatives useful:

$$\begin{aligned} \frac{\partial \text{tr}(\mathbf{WA})}{\partial \mathbf{W}} &= \mathbf{A}^T \\ \frac{\partial \text{tr}(\mathbf{WAW}^T)}{\partial \mathbf{W}} &= \mathbf{WA}^T + \mathbf{WA} \end{aligned}$$

5. (10 points) **Regularized least squares**

In lecture, we worked through the following least squares problem

$$\arg \min_{\theta} \frac{1}{2} \sum_{i=1}^N (y^{(i)} - \theta^T \hat{\mathbf{x}}^{(i)})^2$$

However, the least squares has a tendency to overfit the training data. One common technique used to address the overfitting problem is regularization. In this problem, we work through one of the regularization techniques namely ridge regularization which is also known as the regularized least squares problem. In the regularized least squares we solve the following optimization problem

$$\arg \min_{\theta} \frac{1}{2} \sum_{i=1}^N (y^{(i)} - \theta^T \hat{\mathbf{x}}^{(i)})^2 + \frac{\lambda}{2} \|\theta\|_2^2$$

where λ is a tunable regularization parameter. From the above cost function it can be observed that we are seeking least squares solution with a smaller 2-norm. Derive the solution to the regularized least squares problem, i.e Find θ^* .

6. (20 points) **Linear regression.**

Complete the Jupyter notebook `linear_regression.ipynb`. Print out the Jupyter notebook as a PDF and submit it to Gradescope.