

CSC 390: Topics in Artificial Intelligence

Midterm: Fall 2016

Due:

| |
|------------------------------|
| Wednesday, October 26 at 5pm |
|------------------------------|

- This is a take-home exam with unlimited time from when it is out to when it is due.
- It is open-notes, so you may use any course materials. If you use any online resources that haven't been part of this class, please cite them.
- No communication about the exam with anyone in the class (or outside the class).
- If there is a clarification I think should be made to the entire class, I'll post it on Piazza.
- I will still have office hours Thursday 4-5pm and Monday 5-6pm, but I might not say much!
- If you are not in class to pick up an exam due to Grace Hopper, I will have copies outside my office. You are also welcome to print the exam on your own, it will be available on the course website.
- Turn in your exam to me in person or put it under the door of my office if I am not there (not in the bin outside).
- If you are unable to make it to my office to turn in the exam before the deadline, email me to make other arrangements such as scanning your exam.
- If you are unable to make progress on any part of the exam, tell me what you tried; describe your thought process.

| | |
|------|--|
| Name | |
|------|--|

| | |
|--------|------|
| Part 1 | /25 |
| Part 2 | /15 |
| Part 3 | /25 |
| Part 4 | /20 |
| Part 5 | /15 |
| Total | /100 |

Part 1: Vocab and Written Communication

- (a) Is the nearest-neighbors method an example of supervised or unsupervised learning? Justify your answer.
- (b) Consider the data analysis we did in Homework 4. Was this an example of supervised or unsupervised learning? Justify your answer.
- (c) Is UPGMA an example of agglomerative or divisive clustering? Justify your answer.
- (d) How can labeled *test* data be used to assess the accuracy of *k*-means? Does this qualify as supervised learning? Would the training data need to be labeled to use this approach? Justify your answer.

(e) Consider the training dataset:

$$X = \begin{bmatrix} 2 & -1 & 6 \\ 0 & 3 & -3 \end{bmatrix}$$

with labels $y = [0, 1]$. Given a new test data point $x_{\text{test}} = [-1, 1, 2]$, what label would you assign it using a 1-nearest neighbors approach with Euclidean distance? Show your work to justify your answer.

(f) In k -nearest neighbors we have one parameter, k . In linear regression, we have $p+1$ parameters, where p is the number of features (+1 for the intercept). In these two methods, which parameters are chosen in advance by the user? What data is used to find the parameters *not* chosen by the user? Explain your answer, using the concepts of “fitting a model” and “training data”.

Part 2: Linear Regression

Consider the labeled training dataset $p_1 = (x_1, y_1) = (-1, 1)$ and $p_2 = (x_2, y_2) = (3, 0)$, where each point has one feature and one class label in $\{0, 1\}$.

(a) Plot these two points in 2D, labeling your axes.

(b) After running linear regression on this dataset, what is the resulting linear equation? Show this line on your graph above. How many parameters does this linear model have?

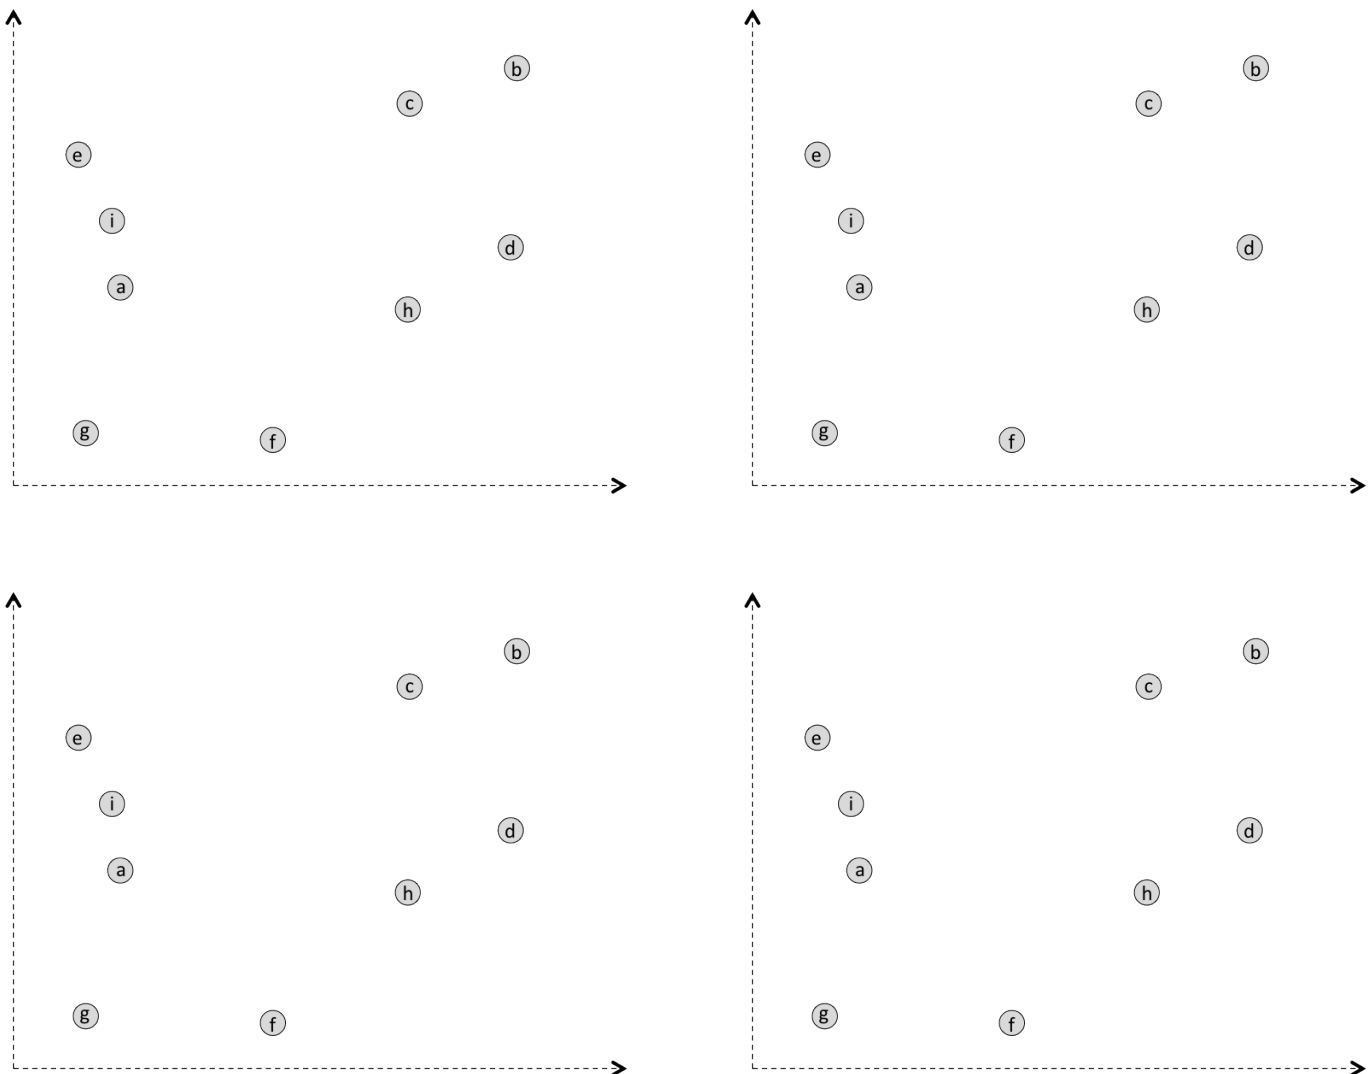
(c) Given a new test data point x_{test} (again with one feature), what rule does this model give us about how x_{test} should be classified?

Part 3: Hierarchical Clustering

Inspired by UPGMA, you decide to create a clustering algorithm that is top-down instead of bottom-up. So instead of first choosing the two points that are the closest to each other, in this new algorithm we will choose the two points that are *furthest away* from each other. This algorithm will proceed as follows. The data points will begin in one cluster. Then iterate the following steps until all points are in their own clusters:

- Select a cluster c_j to split. Find the two points that are the furthest from each other within this cluster. Denote these as x_{m1} and x_{m2} . These two points will now be in different clusters.
- For each point within c_j , compare its distance from x_{m1} and x_{m2} and assign it to the closest point. In this way we break up the cluster c_j into two new clusters: c_{j1} and c_{j2} .

- (a) Given the example dataset in 2D below, run this clustering algorithm. The dataset is replicated a few times to show different iterations.




- (b) Draw out the tree (hierarchy) created by running this clustering algorithm on these points. Can we label the branch lengths using this algorithm? Why or why not?

- (c) What could go wrong when using this algorithm? Sketch an example dataset where this algorithm would not achieve a desirable clustering.

Part 4: Clustering Literature

The figure below shows the β -similarity graph for $\beta = 0.3$, for a set of data points $X = [a, b, c, d, e, f, g, h, i]$. Each value is a measure of similarity between two points (so higher values denote more similar points). Edges are drawn only if the similarity value is at least β .

- (a) Using the β -similarity graph on the left, create the max- S graph on the right, making sure to consider each point in turn.



- (b) Using the *dynamic hierarchical compact* algorithm framework from Paper 2 (i.e. no overlapping clusters), how many clusters (k) should we obtain from the max- S graph? Write out each of these clusters with the points assigned to them. Denote this clustering $\mathcal{C} = \{c_1, c_2, \dots, c_k\}$.

(c) A different clustering algorithm obtains the following clustering $\mathcal{D} = \{d_1, d_2\}$, where

$$d_1 = \{a, d, e, f\} \quad \text{and} \quad d_2 = \{b, c, g, h, i\}.$$

We will compare these two clusterings using the *Rand index*. First, compute A , the number of *pairs* of points that are assigned the *same* cluster in both \mathcal{C} and \mathcal{D} .

(d) Next compute B , the number of pairs of points that are assigned to *different* clusters in *both* \mathcal{C} and \mathcal{D} .

(e) Using these results, compute the Rand index:

$$R = \frac{A + B}{\binom{n}{2}}$$

where n is the number of points ($n = 9$ here). Interpret this quantitative result.

Part 5: Dimensionality Reduction

Consider the following dataset with $m = 3$ data points and $p = 4$ features:

$$X = \begin{bmatrix} -3 & 0 & 2 & 1 \\ 1 & -1 & 0 & 2 \\ 2 & 1 & -2 & -3 \end{bmatrix}$$

Note that the features have been normalized (i.e. each column has mean 0). Without feature selection / dimensionality reduction, we cannot visualize this dataset.

- (a) Below is a portion of the covariance matrix A for this dataset. Fill in the missing entries, showing your work. Will the diagonal always contain non-negative values?

$$A = \begin{bmatrix} & \frac{1}{2} & -5 & -\frac{7}{2} \\ \frac{1}{2} & 1 & & \\ -5 & & 4 & 4 \\ & & 4 & 7 \end{bmatrix}$$

- (b) The eigenvalues of A are: $\lambda_1 = \frac{1}{2}(19 + 4\sqrt{7})$, $\lambda_2 = \frac{1}{2}(19 - 4\sqrt{7})$, $\lambda_3 = 0$, $\lambda_4 = 0$. Which eigenvalue corresponds to the first principal component (PC)? Which eigenvalue corresponds to the second PC? What do the zero eigenvalues tell us about this dataset?

(c) The resulting transformed dataset is approximately:

$$X^* \approx \begin{bmatrix} 4.4 & 6.2 \\ -0.6 & 1.6 \\ -3.8 & -7.8 \end{bmatrix}$$

Plot these points on a graph (approximately), labeling your axes. Then write out the matrix multiplication problem that results in X^* (you don't need to use any numerical values, just show the matrices involved and how the dimensions work out). Using this matrix multiplication set up, explain how each of these two new features is a linear combination of the original features.