

Practice Final Exam

19 December 2012

You have up to 2 hours, 30 minutes. You may not use text book or notes. You may refer to an online spreadsheet in those problems that provide a URL.

1. For each statement below, fill in the blank with the *best* term from the following list. Some terms might be used more than once; some might not be used at all.

- crossover • entropy • Euclidean distance • Manhattan distance • mutation
- sigmoid function • step function • supervised learning • tournament selection
- unsupervised learning

- (a) _____ is a method for calculating distances between points, based on the square root of the sum of the squared differences along each dimension.
- (b) The _____ governs the activation of a (simulated) neuron, where the value jumps from zero to one at some threshold.
- (c) _____ is a way to choose individuals from a population that will probabilistically select those with higher fitness.
- (d) In _____, the learning algorithm is not provided with a correct classification for the data. Instead, it tries to cluster and find patterns.

2. In a genetic algorithm with bit representations, which of the following strings **cannot** be the result of a *crossover* between X and Y?

```
X = 1 1 0 1 0 1 0 0 1 1 1 1 1 0 0 0
Y = 0 1 1 1 0 1 0 1 1 1 0 0 0 1 0 0
```

- (a) 1 1 0 1 0 1 0 0 1 1 1 0 0 1 0 0
- (b) 0 1 1 1 0 1 0 0 1 1 1 0 0 0 1 0
- (c) 1 1 0 1 0 1 0 1 1 1 0 0 0 1 0 0
- (d) 0 1 1 1 0 1 0 0 1 1 1 1 1 0 0 0

3. Which machine learning techniques are most suitable for classifying continuous, real-valued data, such as measurements?

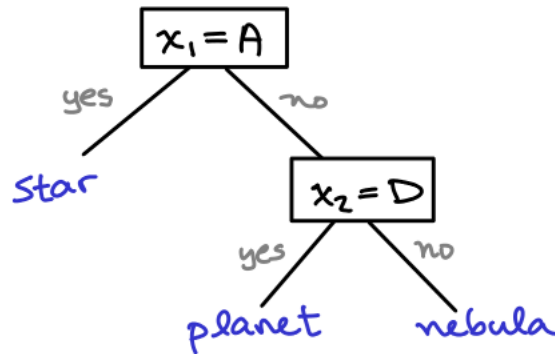
- (a) Decision Tree Induction (ID3)
- (b) Back-Propagation Neural Network
- (c) Genetic Algorithm
- (d) k-Nearest Neighbors

4. Suppose you were hired by the *Space Telescope Science Institute* in Baltimore. This is the organization that collects scientific data from the Hubble Space Telescope. They want you to create an intelligent program that can examine pictures of space objects and classify them as *planets*, *stars*, or *nebulae*. They already have subroutines that can extract four features (named x_1 – x_4) from the images. They provide you with these sample data:

ID	x_1	x_2	x_3	x_4	classification
1	A	D	H	L	star
2	A	E	H	K	star
3	A	F	G	L	planet
4	A	F	H	K	star
5	A	F	H	L	star
6	B	D	G	K	planet
7	B	D	H	K	nebula
8	B	E	G	K	planet
9	C	D	G	L	planet
10	C	E	G	L	planet
11	C	E	H	L	nebula

You can view these examples in a Google Spreadsheet at <http://bit.ly/162galaxy> — I suggest trying View » List on the menu.

- (a) Below is a colleague's attempt to design a decision tree for these data. Which records does it misclassify?

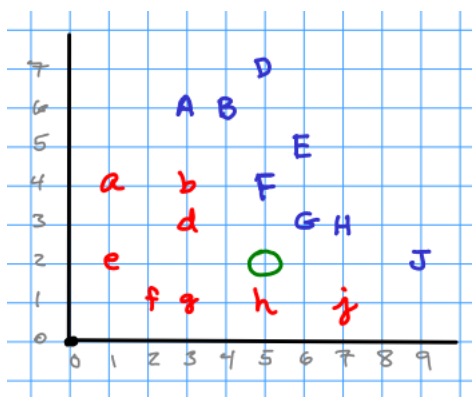


(b) Use the telescope data and my entropy calculator at <http://bit.ly/162entropy> to determine the *best* initial test, and then produce a decision tree that will correctly classify all 11 records.

(c) Using the decision tree you developed in the previous part, classify these two additional records:

ID	x_1	x_2	x_3	x_4	classification
12	A	D	G	L	_____
13	B	F	H	L	_____


5. Below is a Cartesian coordinate system containing 8 points classified as blue (upper-case letters), and 8 points classified as red (lower-case letters).



Your task is to classify the green circle at (5, 2) using a K *nearest-neighbors* algorithm. To keep the numbers simple, use Manhattan (city-block) distance.

- What classification (red/blue) will be chosen when $K = 1$ (meaning the algorithm considers just the single nearest neighbor).
- What classification will be chosen when $K = 3$ (meaning the three nearest neighbors *vote* on the classification).
- What classification will be chosen when $K = 7$?

6. Below is a multi-layer perceptron with three input neurons, two hidden neurons, and one output. Each neuron takes the sum of its weighted inputs, and applies a step function (as shown) to determine its own activation.



- (a) What is the output activation when the three inputs (left to right) are 0 1 0?
- (b) What is the output activation when the three inputs (left to right) are 1 0 1?
7. In class and on assignment 6, we used a genetic algorithm to solve a *knapsack problem*, where we decide which objects to pack into a finite container to maximize the value. That was a 0-1 knapsack, where you either take an object or not. A generalization of this problem considers taking some discrete number of objects of a particular type.

item	count	value(\$)	weight(kg)
A	6	2	3
B	2	3	5
C	3	4	6

In this example, we can pack **up to 6** instances of object A (worth \$2 each, and weighing 3kg each), and so on for B, C. However, our container has a weight limit of 25kg.

Use 7 bits to represent solutions to this problem: 3 bits for the number of 'A' objects, 2 bits for 'B', and 2 bits for 'C'. We'll define fitness as the monetary value, or zero if the selection is over the weight limit, or uses more than the number of available objects in each category.

What is the fitness of each bit string below?

- (a) 0 0 1 0 1 1 1
 (b) 0 1 0 0 0 1 0
 (c) 0 1 1 1 1 0 1
 (d) 1 0 0 0 1 0 1