CS-171, Intro to A.I. — Quiz#5 — Spring Quarter, 2011 — 30 minutes

YOUR NAME AND EMAIL ADDRESS: _____

YOUR ID: _____ ID TO RIGHT:_____ ROW:_____ NO. FROM RIGHT:_____

1. (45 pts total, 3 pts each) Machine Learning concepts.

For each of the following items on the left, write in the letter corresponding to the best answer or the correct definition on the right. The first one is done for you as an example.

A .	Learning	A	Improves performance of future tasks after observing the world
J	Information Gain	B	Fixed set, list, or vector of features/attributes paired with a value
M	Decision Boundary	C	Agent learns patterns in the input with no explicit feedback
P	Support Vector Machine	D	Agent observes input-output pairs & learns to map input to output
L	Cross-validation	E	Example input-output pairs, from which to discover a hypothesis
N	Linear Classifier	F	Examples distinct from training set, used to estimate accuracy
B	Factored Representation	G	Supervised learning with a discrete set of possible output values
D	Supervised Learning	H	Supervised learning with numeric output values
F	Test Set	I	Internal nodes test a value of an attribute, leaf nodes=class labels
O	Naïve Bayes Classifier	J	Expected reduction in entropy from testing an attribute value
G	Classification	K	Choose an over-complex model based on irrelevant data patterns
I	Decision Tree	L	Randomly split the data into a training set and a test set
H	Regression	M	Surface in a high-dimensional space that separates the classes
E	Training Set	N	Tests **w**·**f**>0, where **w** is a weight vector and **f** is a feature vector
C	Unsupervised Learning	O	Tests P (C) Π_i P(X_i | C), where C is a class label and X_i are features
K	Overfitting	P	Current most-popular "off-the-shelf" supervised learning method

2. (25 pts total) Decision Tree Learning.

You are given the following set of animal examples, each of which is assigned a class label of + or — depending on whether or not it is a member of the target class:

Example	Environment?	Mammal?	Class
Ant	Land	No	—
Bat	Land	Yes	+
Dolphin	Sea	Yes	—
Lizard	Land	No	—
Sea Lion	Sea	Yes	—
Zebra	Land	Yes	+
Shark	Sea	No	—
Mouse	Land	Yes	+
Chicken	Land	No	—

	Environment?	Mammal?
	Land/Sea	Y/N
+	3/0	3/0
—	3/3	2/4

2a. (15 pts) Draw the decision tree that would be constructed by recursively applying information gain to select roots of sub-trees, as in the Decision-Tree-Learning algorithm.

Half credit for the correct root; half credit for wrong root but correct classification; full credit for the correct tree.

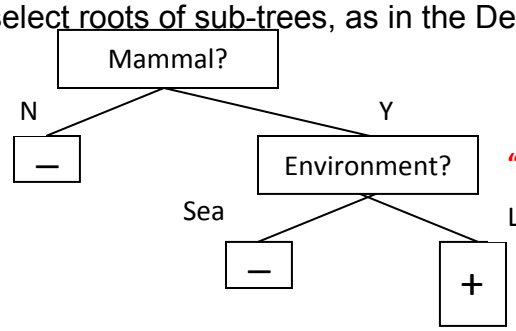

"Land? Y/N" or "Sea? Y/N" here is also OK

2b. (5 pts) What class is Tiger? (Environment=Land, Mammal=Yes)_____+_____

2c. (5 pt) What class is Tuna? (Environment=Sea, Mammal=No)___—____

****** TURN PAGE OVER. QUIZ CONTINUES ON THE REVERSE. ******

4. (30 pts total, 6 pts each) Linear Classifier (Perceptron) Learning.
Label the following as Y (= yes) or N (= no) depending on whether a linear classifier with a "hard" decision boundary (= a perceptron with a step transfer function) can correctly classify the examples shown. If your answer is Y (= yes), fill in a set of weights that correctly classifies them. Use w0 as the threshold and wi as the weight for input xi.

All perceptrons have three Boolean inputs, x1, x2, and x3, and a "dummy" input, x0, which is always equal to one. They all compute the decision function $\sum w_i x_i > 0$. You may not transform the input space, i.e., they operate on the stated inputs.

Recall that the three Boolean inputs map each possible example to a corner of a three-dimensional cube (i.e., to the vector space within which each of your examples is a point). Can a linear classifier decision boundary represent the named concept?

The first one is done for you as an example.

4a. (3 pts) "At least two inputs are 1."
List the named set S, represented as 3-bit Boolean numbers: S={011,101,110,111} .
Correctly classifiable? Y
If yes, weights are w0 = -1.5 ; w1 = 1 ; w2 = 1 ; w3 = 1

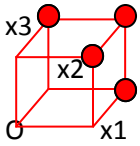

4b. (3 pts) "Exactly two inputs are 1."
List the named set S, represented as 3-bit Boolean numbers: S={011,101,110} .
Correctly classifiable? N
If yes, weights are w0 = ___; w1 = _____; w2 = _____; w3 =_____

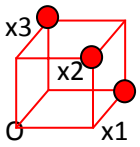

4c. (3 pts) "At most two inputs are 1."
List the named set S, represented as 3-bit Boolean numbers:
S={000,001,010,100,011,101,110} .
Correctly classifiable? Y
If yes, weights are w0 = 2.5 ; w1 = -1 ; w2 = -1 ; w3 = -1

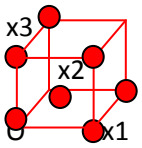

4d. (3 pts) "Input x1 = 1, input x2 = 0, input x3 =anything."
List the named set S, represented as 3-bit Boolean numbers: S={100,101} .
Correctly classifiable? Y
If yes, weights are w0 = -0.5 ; w1 = 1 ; w2 = -1 ; w3 = 0

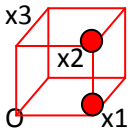

4e. (3 pts) "IF input x1 = 1 THEN input x2 = 0 ELSE input x2 = 1."
List the named set S, represented as 3-bit Boolean numbers: S={010,011,100,101} .
Correctly classifiable? N
If yes, weights are w0 = _____; w1 = _____; w2 = _____; w3 =_____

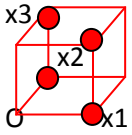

4f. (3 pts) "Input x1 = input x2."
List the named set S, represented as 3-bit Boolean numbers: S={000,001,110,111} .
Correctly classifiable? N
If yes, weights are w0 = _____; w1 = _____; w2 = _____; w3 =_____