CS-171, Intro to A.I. — Quiz#4 — Winter Quarter, 2012 — 20 minutes

YOUR NAME AND EMAIL ADDRESS: _____

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

1. (30 pts total, 2 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 (Feature Vector)	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 $P(C) \prod_i P(X_i \mid C)$, where C is a class label and X_i are features
K	Overfitting	P	Current most-popular "off-the-shelf" supervised learning method

2. (30 pts total, 10 pts each) Bayesian networks.
For each Bayesian network shown below, write down in factored form the joint probability distribution that it represents.

2a. P(C | A, B) P(A) P(B)

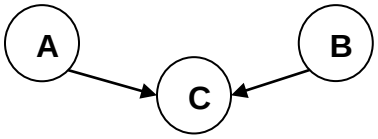

2b. P(A | C) P(B | C) P(C)

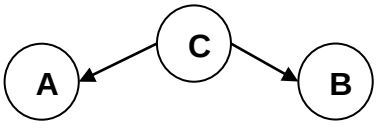

2c. P(C | B) P(B | A) P(A)

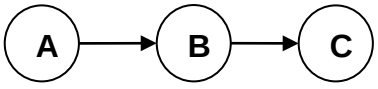

****** TURN PAGE OVER. QUIZ CONTINUES ON THE REVERSE ******

2. (25 pts total) Decision Tree Learning.
You are an agricultural robot given the following set of plant examples. Each is assigned a class label of + or — depending on whether or not it is a member of the target class:

Example	Vine?	Fruit?	Leaf?	Class
Watermelon	Yes	Yes	Curly	+
Ivy	Yes	No	Curly	—
Bougainvillea	Yes	No	Flat	—
Kudzu	Yes	No	Flat	—
Maple	No	No	Curly	+
Oak	No	No	Flat	+
Sycamore	No	No	Flat	+
Apple	No	Yes	Curly	—

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.

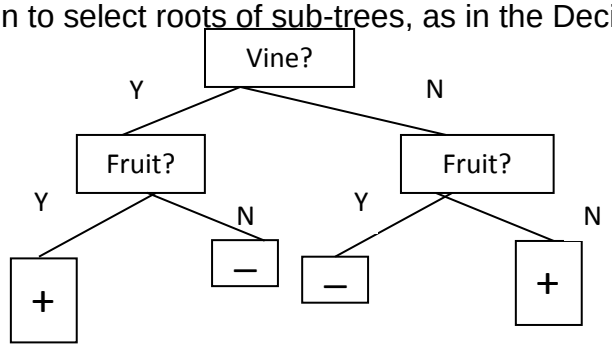

Full credit if your answers are right for the tree you drew, even if the tree itself is wrong.

2b. (5 pts) What class is Grape? (Vine=Yes, Fruit=Yes, Leaf=Curly) +

2c. (5 pt) What class is Orange? (Vine=No, Fruit=Yes, Leaf=Curly) —

4. (10 pts total) Bayesian networks.
Draw the Bayesian network that represents P(J | A) P(M | A) P(A | B, E) P(B) P(E).