

1	2	3	4	5	6	7	total
20	20	20	30	20	20	25	155

UMBC CMSC 471 01, Final Exam, 21 May 2018

Name: _____

Please write all of your answers on this exam. The exam is closed book and has seven problems that add up to 155 points. You have the two hours to work on this exam. Good luck.

1. True/False (20 points) Circle T or F for each statement

- T F Every complete inference procedure is also sound. **True**
- T F Every sound inference procedure is also complete. **False**
- T F A sound but incomplete inference procedure can produce an incorrect answer. **False**
- T F An unsound inference procedure can produce an incorrect answer. **True**
- T F Every model of $(a \wedge (b \vee c))$ is also a model of $(a \vee (b \vee c))$. **True**
- T F If a propositional sentence is valid, it is satisfied in all possible models. **True**
- T F The STRIPS planning algorithm was designed to find the shortest possible plan for a given problem. **False**
- T F Regression planners search for plans by starting from a goal state and working back toward the initial state. **True**
- T F Decision tree learning can only be used to produce binary classifiers. **False**
- T F In a well-formed Bayesian Belief Network, a node is always conditionally independent of its non-descendants given its parents. **True**
- T F Information gain is used to determine the network structure in a Bayesian Network. **False**
- T F Bayes' rule can be used to relate the probability of a cause given its symptoms to the probability of its symptoms given a cause. **True**
- T F An important advantage of support vector machines (SVMs) is that they can directly implement classifiers with a large number of classes. **False**
- T F The K-nearest neighbors algorithm is an unsupervised machine learning technique for clustering instances into K clusters where K is a positive, non-zero integer. **False**
- T F One drawback of the K-means algorithm is that one needs to specify exactly how many clusters the algorithms should find. **True**
- T F Bagging is a technique used in building Bayesian Belief Networks that helps identify variables that are not statistically independent. **False**
- T F Overfitting occurs when a machine learning model was trained with two examples with different labels. **False**
- T F The precision metric used in machine learning is defined as the ratio of the number of true positives to the sum of the number of true and false positive. **True**
- T F The precision metric used in machine learning is defined as the ratio of the number of true positives to the sum of the number of true and false positive. **True**
- T F A learning curve in machine learning shows the trade-off between a system's precision and recall. **False**

2. Multiple choice (20 points: 5, 5, 5, 5)

Circle the letters of **all** of the correct answers for each question.

2.1 Bayesian Belief networks ...

- (a) Have an associated directed, acyclic graphs **True**
- (b) Encode conditional independence assertions among random variables **True**
- (c) Generally require less storage than the full joint distribution **True**
- (d) Make the assumption that all parents of a single child are independent given the child
- (e) None of the above

2.2 How many models does the propositional sentence $A \vee (B \Rightarrow C)$ have?

- (a) one
- (b) three
- (c) five
- (d) seven **True**
- (e) eight
- (f) None of the above

2.3 What tasks are appropriate for unsupervised machine learning algorithms like clustering?

- (a) predicting if a credit card transaction is fraudulent or legitimate
- (b) finding items that are frequently bought together in a convenience store **True**
- (c) organizing a person's connections in a social media network into groups **True**
- (d) estimating how much a house will cost given its square footage
- (e) none of the above

2.4 What are potential advantages of using a random forest classifier over a decision tree classifier?

- (a) It tends to give more accurate results **True**
- (b) It can handle a much larger number of features than a simple decision tree
- (c) It can make decisions based on non-linear combinations of features
- (d) It is less prone to overfitting the training data **True**
- (e) none of the above

3. Short Answers (20 points: 1,; 10)

3.1 In a few sentences, describe the approach used in the classic STRIPS planning algorithm.

The STRIPS planner compared the current and goal states to find a list of facts that are true in the goal state but not the current one. It then loops over this list of missing facts and, for each, looks for an action that has the missing fact as one of its effects. If one is found, the planner tries to execute the action using a recursive call to the planner in which the actions preconditions are a new goal state. If that process succeeds, the action is performed, and the loop continues. If it fails, then another action that has the missing action as one of its effects is tried. When all of the condition in the goal state have been achieved the planner stops with success.

3.2 Describe one problem that the simple STRIPS planning algorithm has that can cause it to produce a plan that is longer than necessary.

The planner is prone to the “Sussman anomaly”, in which it achieves a necessary condition in an initial part of the plan only to undo the condition later in the plan in order to accomplish some other goal. Eventually, the planner will note that the first condition has become unsatisfied and achieve it again. But the result is a plan with needless steps.

4. Resolution Proof in Propositional Logic (30: 10, 5, 15)

Consider the following propositional variables:

R - it is raining

U - you have an umbrella

W - you get wet

and the following statements, which can be expressed as implications:

If it rains, you always have an umbrella ($R \Rightarrow U$)

If you have an umbrella, you do not get wet ($U \Rightarrow \sim W$)

If it is not raining, you do not get wet ($\sim R \Rightarrow \sim W$)

A propositional sentence is in CNF if it is a set of one or more expressions where each is a disjunction of variables or negated variables

4.1 Construct a KB of propositional sentences using the three propositional variables (R, U, W) and logical connectives (\wedge , \vee , \neg , \rightarrow). Encode each of the three English statements into *one or more* logic sentences in conjunctive normal form (CNF). (10 pts)

#	English statements	CNF clauses (one or more)
0	If it rains, you always have an umbrella	0.1 $\sim R \vee U$
		0.2
		0.3
1	If you have an umbrella, you do not get wet	1.1 $\sim U \vee \sim W$
		1.2
		1.3
2	If it is not raining, you do not get wet	2.1 $R \vee \sim W$
		2.2
		2.3

4.2 How many models are there for this KB? Recall that a model is an assignment of true and false to each variable such that the KB is consistent. (5 points)

3

R	U	W
1	1	0
0	1	0
0	0	0

4.3 Show a resolution refutation proof for “you will not get wet”. That is, given your KB, prove that you will not get wet. Start with the negation of what’s to be proved, add statements from your KB and then use resolution to derive a contradiction (\perp). The table to the right shows an example resolution refutation proof (15 points).

step	action	result
1	assume	$\neg Q$
2	KB	$\neg P \vee Q$
3	KB	P
4	resolve 2,3	Q
5	resolve 1,4	\perp

Sample proof of Q given $P \rightarrow Q$ & P

step	action	result
1	assume	W
2	KB	$\sim R \vee U$
3	KB	$\sim U \vee \sim W$
4	KB	$R \vee \sim W$
5	Resolve 2, 4	$U \vee \sim W$
6	Resolve 3, 5	$\sim W$
7	Resolve 1, 6	\perp
8		
9		
10		
11		

5. English and logic (20 points, 5, 5, 5, 5)

For each English sentence below, write the **letter or letters** corresponding to **all** of the logic expressions that express its meaning. If none of the logic expressions correspond to the English sentence, write **none** in the box.

5.1 Every person is either male or female

B

- A. $\forall x \text{ person}(x) \Rightarrow \text{male}(x) \wedge \text{female}(x)$
- B. $\sim \exists x \text{ Person}(x) \wedge \sim (\text{male}(x) \vee \text{female}(x))$
- C. $\forall x (\text{Person}(x) \Rightarrow \text{male}(x)) \wedge (\text{Person}(x) \Rightarrow \text{female}(x))$
- D. $\exists x \text{ person}(x) \Rightarrow \text{male}(x) \wedge \text{female}(x)$

5.2 No person can be both male and female

D

- A. $\forall x \text{ person}(x) \Rightarrow \text{male}(x) \vee \text{female}(x)$
- B. $\sim \exists x \text{ person}(x) \wedge \sim (\text{male}(x) \vee \text{female}(x))$
- C. $\forall x (\text{person}(x) \Rightarrow \text{male}(x)) \vee (\text{person}(x) \Rightarrow \text{female}(x))$
- D. $\sim \exists x \text{ person}(x) \wedge \text{male}(x) \wedge \text{female}(x)$

5.3 Every person has a mother

C

- A. $\forall x \exists y \text{ person}(x) \wedge \text{hasMother}(x, y)$
- B. $\exists y \forall x \text{ person}(x) \Rightarrow \text{hasMother}(x, y)$
- C. $\forall x \exists y \text{ person}(x) \Rightarrow \text{hasMother}(x, y)$
- D. $\exists x \exists y \text{ person}(x) \Rightarrow \text{hasMother}(x, y)$

5.4 Good food is not cheap and cheap food is not good

D

- A. $\forall x \text{ food}(x) \Rightarrow (\text{good}(x) \wedge \sim \text{cheap}(x)) \vee (\sim \text{good}(x) \wedge \text{cheap}(x))$
- B. $\forall x \text{ food}(x) \Rightarrow (\text{good}(x) \Leftrightarrow \sim \text{cheap}(x))$
- C. $\forall x \text{ food}(x) \Rightarrow (\sim \text{good}(x) \Leftrightarrow \text{cheap}(x))$
- D. $\text{food}(x) \Rightarrow (\text{good}(x) \Rightarrow \sim \text{cheap}(x)) \wedge (\text{cheap}(x) \Rightarrow \sim \text{good}(x))$

6. Probabilistic Reasoning (20 pts: 5;5;5;5)

The table at the right gives the full joint probability distribution for Boolean variables A, B and C. Calculate the following probabilities. Show your work for partial credit.

A	B	C	p(A,B,C)
t	t	t	0.03
t	t	f	0.12
t	f	t	0.17
t	f	f	0.18
f	t	t	0.03
f	t	f	0.12
f	f	t	0.24
f	f	f	0.11

6.1 $p(A=f) = 0.5$

6.2 $p(B=t) = 0.3$

6.3 $p(B=t, C=t) = 0.06$

6.4 $p(A=f, C=t) = 0.27$

6.5 $p(A=t | B=t) = 0.5$

6.6 $p(C=f | B=t) = 0.8$

6.7 Are A and B independent of each other? (Y=Yes, N=No): **Y**

6.8 Are B and C independent of each other? (Y=Yes, N=No): **N**

6.9 Are B and C conditionally independent given A? (Y=Yes, N=No): **N**

6.10 Are A and C conditionally independent given B? (Y=Yes, N=No): **Y**

Recall that

$$P(a, b) = P(a|b) * P(b) \\ = P(b|a) * P(a)$$

so

$$P(a|b) = P(b|a) * P(a)/P(b)$$

7. Decision tree classification (25)

NASA has secretly collected aliens and identified them as coming from either Mars (M) or Venus (V). You are charged with building a decision tree classifier to predict the source of new aliens (M or V) as they are found using four predictive features: height, color, number of legs and stench.

<i>predictive features</i>				<i>target</i>
height	color	# legs	stench?	source
short	green	3	yes	M
tall	green	3	no	M
short	green	3	no	M
tall	green	3	no	M
tall	red	2	yes	M
tall	green	2	yes	V
short	red	2	no	V
tall	red	3	no	V
short	red	3	no	V
tall	red	3	no	V

7.1 What is the initial entropy of the target variable source, i.e. $E(\text{source})$, which is defined as the sum of the probability of each value times the negative of the log base 2 of that probability. (5)

$$E(S) = \sum_{i=1}^c -p_i \log_2 p_i$$

1.0

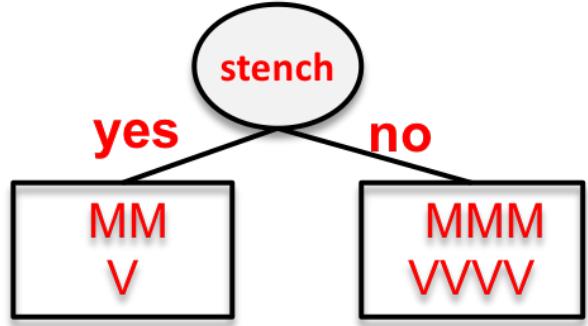
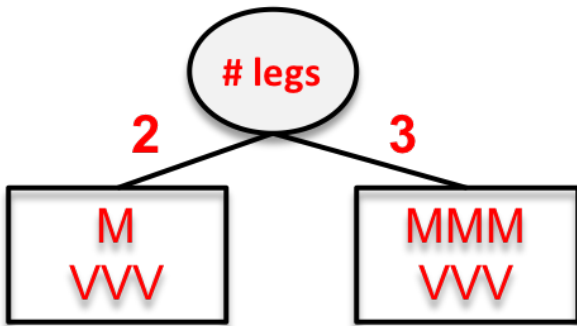


p	log₂(p)
0.5	-1.0

7.2 Fill in the diagrams on the next page that show the four possibilities of an initial split on each of the four variables. For each one, (1) label the arcs from the root node to the boxes and (2) indicate the source values that go in each box. We've done the first one for you. (5)

7.3 Circle the variable that an ID3-type decision tree algorithm will choose for the initial split. (5)

Height Color # Legs Stench

7.4 Using your answer in 7.3, draw **the entire** decision tree on the next page that an ID3-type algorithm that uses information gain to choose splits will produce. (10)



Complete the diagram above for question 7.3 and draw your complete decision tree below

