

CS-171, Intro to A.I. — Final Exam — Fall Quarter, 2015

YOUR NAME: _____

YOUR ID: _____ ID TO RIGHT: _____ ROW: _____ SEAT: _____

The exam will begin on the next page. Please, do not turn the page until told.

When you are told to begin the exam, please check first to make sure that you have all eleven pages, as numbered 1-11 in the bottom-right corner of each page. We wish to avoid copy problems. We will supply a new exam for any copy problems.

The exam is closed-notes, closed-book. No calculators, cell phones, electronics.

Please turn off all cell phones now. No electronics are allowed at any point of the exam.

Please clear your desk entirely, except for pen, pencil, eraser, a blank piece of paper (for scratch pad use), and an optional water bottle. Please write your name and ID# on the blank piece of paper and turn it in with your exam.

This page summarizes the points for each question, so you can plan your time.

1. (15 pts total, 3 pts each) Linear Regression.
2. (4 pts total, 1 pt each) Task Environment.
3. (18 pts total) Clustering.
4. (16 pts total) Constraint Satisfaction Problems and Job Shop Scheduling.
5. (10 pts total) ONE FISH, TWO FISH, RED FISH, BLUE FISH. Resolution Theorem Proving.
6. (15 pts total, 3 pts each) English to FOL Conversion.
7. (12 pts total, 4 pts each) Bayesian Networks.
8. (10 pts total, -1 pt each error, but not negative) Search Properties.

The Exam is printed on both sides to save trees! Work both sides of each page!

1. (15 pts total, 3 pts each) Linear Regression.

Suppose we have following dataset for a linear regression problem.

Training data

Example	Feature $X1$	Feature $X2$	Target Y
Example #1	3	5	18
Example #2	7	2	28

Testing data

Example	Feature $X1$	Feature $X2$	Target Y
Example #3	4	1	18

There are two features in this problem and so we have three parameters (weights) in this regression problem. Parameter θ_0 is the constant weight. Parameters θ_1 and θ_2 are the weights of features $X1$ and $X2$ respectively. Recall that the resulting linear regression model is $\hat{Y} = \theta_0 + \theta_1 * X1 + \theta_2 * X2$.

For this problem, assume that you are given **Weights** = [$\theta_0 = 1, \theta_1 = 3, \theta_2 = 2$].

1.a (9 pts total, 3 pts each) Given **Weights** above, what is the predicted value for the three examples above? **SHOW YOUR WORK.** First display the formula above for \hat{Y} after substituting numbers for θ_i and X_i (this step is worth most or all of the points). Then simplify to obtain a numerical answer. The numbers have been chosen to be so simple that you do not need a calculator.

Predicted value for Example #1:

$\hat{Y}_1 =$ _____

Predicted value for Example #2:

$\hat{Y}_2 =$ _____

Predicted value for Example #3:

$\hat{Y}_3 =$ _____

1.b (6 pts total, 3 pts each) Calculate the value of Mean Squared Error (MSE) for training and testing data separately. Recall that the MSE formula is $MSE = (1/m) \sum_j (Y_j - \hat{Y}_j)^2$, where m is the number of examples.

SHOW YOUR WORK. First display the formula above for MSE after expanding the summation and substituting numbers for $m, Y_j,$ and \hat{Y}_j (this step is worth most or all of the points). Then simplify to obtain a numerical answer. The numbers have been chosen to be so simple that you do not need a calculator. Your answer to **1.b** is correct if it is correct relative to your answer to **1.a**, even if your answer to **1.a** was wrong.

Mean Squared Training Error:

MSE = _____

Mean Squared Testing Error:

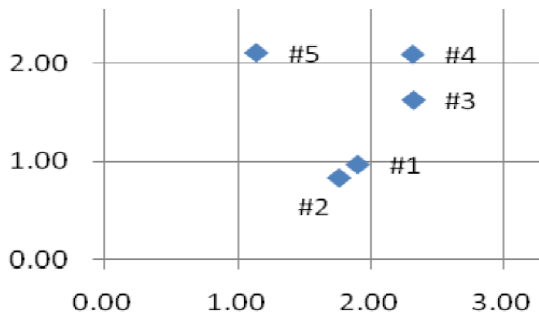
MSE = _____

**** TURN PAGE OVER AND CONTINUE ON THE OTHER SIDE ****

2. (4 pts total, 1 pt each) Task Environment. Your book defines a task environment as a set of four things, with acronym PEAS. Fill in the blanks with the names of the PEAS components.

P _____ E _____ A _____ S _____

3. (18 pts total) Clustering. In this problem you will cluster the following data points, first by K-Means Clustering (3.a) and then by Hierarchical Agglomerative Clustering (3.b). Use Euclidean distance.



Data Point	x	y
Example #1	1.90	0.97
Example #2	1.76	0.84
Example #3	2.32	1.63
Example #4	2.31	2.09
Example #5	1.14	2.11

Distances	#1	#2	#3	#4	#5
#1	-	0.191	0.782	1.193	1.370
#2	-	-	0.968	1.366	1.413
#3	-	-	-	0.460	1.274
#4	-	-	-	-	1.170

3.a. (8 pts total, 2 pts each) K-Means Clustering. Assume the number of clusters is fixed at TWO ($K = 2$). Assume that the initial random assignment of cluster means (= cluster centers) assigns ClusterMean1 = Example #1 and ClusterMean2 = Example #2. Recall that each example is assigned to its nearest cluster mean (= cluster center).

3.a.1 (4 pts total, 2 pts each) Write the example #s that initially are assigned to those clusters, assuming initial cluster means as above. Write your answer as a list of examples in the #i format as given above, i.e., as: #i, #j, #k, ...

(Initial ClusterMean1 = Example #1)

Initial Cluster1 Members = _____

(Initial ClusterMean2 = Example #2)

Initial Cluster2 Members = _____

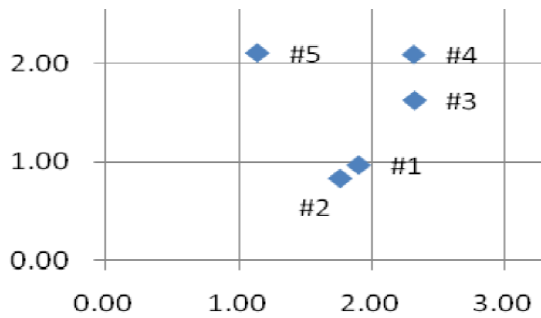
3.a.2 (4 pts total, 2 pts each) Hand-simulate the K-means algorithm until it converges (this happens on the next iteration). The numbers have been chosen to be so simple that you do not need a calculator.

Write the final example #s found in those clusters when K-means converges. Write your answer as a list of examples in the #i format as given above, i.e., as: #i, #j, #k, ...

Final Cluster1 Members = _____

Final Cluster2 Members = _____

3.b. (10 pts total) Hierarchical Agglomerative Clustering. Problem 3.b uses the same data set as problem 3.a.

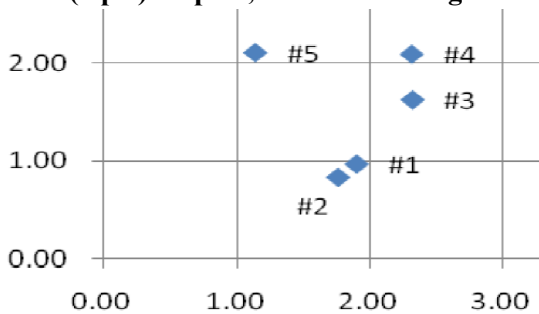


Data Point	x	y
Example #1	1.90	0.97
Example #2	1.76	0.84
Example #3	2.32	1.63
Example #4	2.31	2.09
Example #5	1.14	2.11

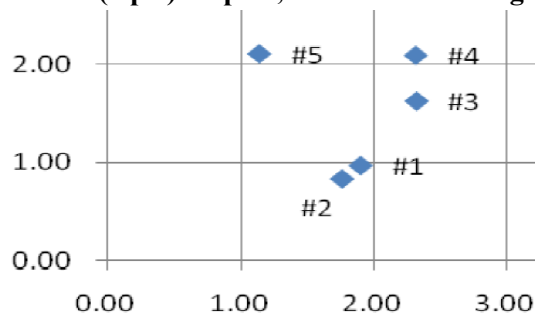
Distances	#1	#2	#3	#4	#5
#1	-	0.191	0.782	1.193	1.370
#2	-	-	0.968	1.366	1.413
#3	-	-	-	0.460	1.274
#4	-	-	-	-	1.170

3.b.1-4 (2 pts each) Hand-simulate hierarchical agglomerative clustering until it converges (this happens after four steps). **Use Min Distance for Cluster Distance**, i.e., distance between two clusters is the minimum distance from any point in one cluster to any point in the other cluster. Initially, each point is one cluster. Then, at each step, the two closest clusters are merged to create a new cluster. **At each step, circle the points in the new merged cluster formed at that step.**

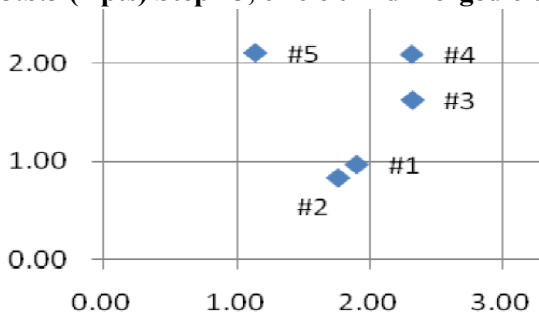
3.b.1 (2 pts) Step #1, circle first merged cluster.



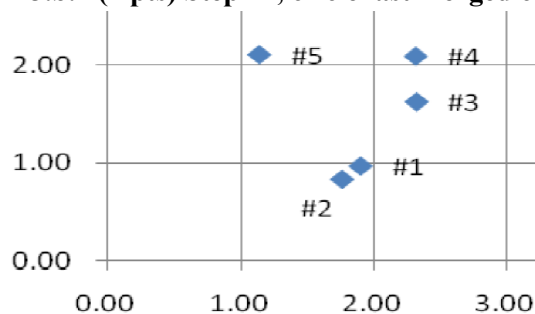
3.b.2 (2 pts) Step #2, circle second merged cluster.



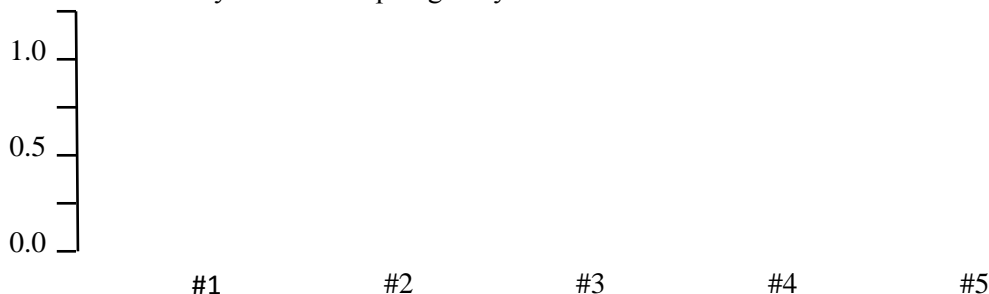
3.b.3 (2 pts) Step #3, circle third merged cluster.



3.b.4 (2 pts) Step #4, circle last merged cluster.



3.b.5 (2 pts) Draw the cluster dendrogram that results. The x-axis gives the example #s. The y-axis gives the distance between the merged clusters at which the new cluster was formed. It is OK to give only very approximate y values, i.e., you will receive full credit if your tree is topologically correct.



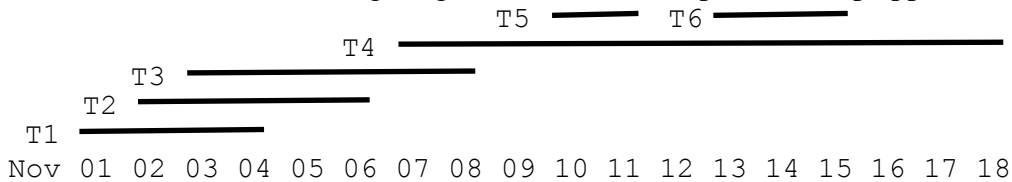
**** TURN PAGE OVER AND CONTINUE ON THE OTHER SIDE ****

4. (16 pts total) **Constraint Satisfaction Problems and Job Shop Scheduling.** You are a job shop scheduling robot that must assign workers to tasks in order to meet scheduling constraints. You have a list of tasks that must be done, a list of workers who can do each task, and a list of times at which each task must be done. (In general, such planning problems are more complicated than our constraint satisfaction problems because their constraints can change as time passes, but this problem has been simplified to match the material you studied.)

Your list of 6 tasks and when they must be done is:

- Task T1: must be done from Nov 1 to Nov 4. It will occupy one worker for all 4 days.
- Task T2: must be done from Nov 2 to Nov 6. It will occupy one worker for all 5 days.
- Task T3: must be done from Nov 3 to Nov 8. It will occupy one worker for all 6 days.
- Task T4: must be done from Nov 7 to Nov 18. It will occupy one worker for all 12 days.
- Task T5: must be done from Nov 10 to Nov 11. It will occupy one worker for all 2 days.
- Task T6: must be done from Nov 13 to Nov 15. It will occupy one worker for all 3 days.

For your convenience, a timeline giving tasks and their temporal overlap appears below:



Your list of 3 workers and the tasks they can do is:

- Worker W1: can do tasks T3, T4, T5, and T6.
- Worker W2: can do tasks T1, T2, T4, T5, and T6.
- Worker W3: can do tasks T1, T3, T5 and T6.

Your scheduling constraints are:

- Every task must be started on its assigned start date.
- Every task must continue for the length of time stated above for that task.
- Every task must be assigned to a worker who can do that task.
- Once a worker is assigned to a task that worker must work only on that task until finished.
- Once a worker is assigned to a task that worker cannot be assigned to another task until finished.
- Only one worker may be assigned to each task.
- Therefore, tasks that overlap in time cannot be assigned to the same worker.

You are overwhelmed by the complexity of your assignment. So, you decide to formulate your assignment as a Constraint Satisfaction Problem. You create **one variable for each task**, and so you have six variables, T1-T6. The **domain of each variable is the set of workers who can do that task**. The **constraints are that every pair of tasks that overlap in time must have a different worker** for each such task (similar to the AllDiff constraint in Sudoku problems). Solving the resulting Constraint Satisfaction Problem yields your job schedule.

4.a (6 pts total, 1 pt each). **Enumerate the domains of each variable.** For each task variable T1-T6 named below, list the domain values of the workers W1-W3 who can do that task.

4.a.1 (1 pt) Task T1: domain = _____

4.a.2 (1 pt) Task T2: domain = _____

4.a.3 (1 pt) Task T3: domain = _____

4.a.4 (1 pt) Task T4: domain = _____

4.a.5 (1 pt) Task T5: domain = _____

4.a.6 (1 pt) Task T6: domain = _____

4.b (6 pts total, 1 pt each). Enumerate the constraints among overlapping variables. For each scope (= pair of task variables T1-T6 that overlap in the timeline listed above), **list the consistent pairs of variable values**, i.e., those pairs that are allowed by the constraints. One constraint is done for you as an example.

4.b.1 Scope = (T4, T6). Consistent pairs = (W1, W2), (W1, W3), (W2, W1), (W2, W3) _____

4.b.2 (1 pt) Scope = (T1, T2). Consistent pairs = _____

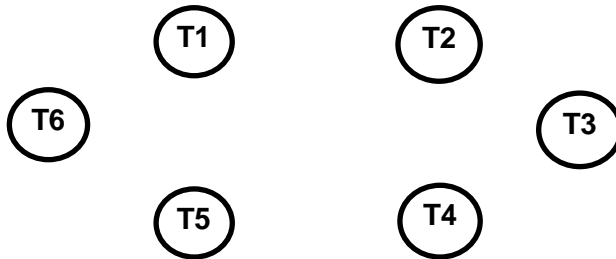
4.b.3 (1 pt) Scope = (T1, T3). Consistent pairs = _____

4.b.4 (1 pt) Scope = (T2, T3). Consistent pairs = _____

4.b.5 (1 pt) Scope = (T3, T4). Consistent pairs = _____

4.b.6 (1 pt) Scope = (T4, T5). Consistent pairs = _____

4.c (1 pt total) Draw the constraint graph for variables T1-T6.



4.d (4 pts total, 2 pts each) Variable ordering during Backtracking Search. Assume that T2 is already assigned to its only domain value. Assume that all other variables are unassigned and have the domain values listed in (4.a) above. Assume that no constraint propagation has been done.

4.d.1 (2 pts) What variables might be assigned by the Minimum Remaining Values (MRV) heuristic:

4.d.2 (2 pts) What variables might be assigned by the Degree Heuristic (DH) heuristic:

**** TURN PAGE OVER AND CONTINUE ON THE OTHER SIDE ****

5. (10 pts total) ONE FISH, TWO FISH, RED FISH, BLUE FISH. Resolution Theorem Proving. (With apologies to Dr. Seuss.)

Amy, Betty, Cindy, and Diane went out to lunch at a seafood restaurant. Each ordered one fish. Each fish was either a red fish or a blue fish.

Their waiter reported:

**“Amy and Betty had fish of different colors; I don’t remember who had what.
 Betty and Cindy had fish of different colors; I don’t remember who had what.
 Cindy and Diane had fish of different colors; I don’t remember who had what.
 Amy, Cindy, and Diane had exactly two red fish among them; I don’t remember who had what.”**

You translate these facts into Propositional Logic (in infix form) as:

/* Ontology: Symbol A/B/C/D means that Amy/Betty/Cindy/Diane had a red fish. */

$(A \iff (\neg B)) \quad (B \iff (\neg C)) \quad (C \iff (\neg D))$
 $((A \wedge C \wedge (\neg D)) \vee (A \wedge (\neg C) \wedge D) \vee ((\neg A) \wedge C \wedge D))$

Betty’s daughter asked, “Is it true that my mother had a blue fish?”

You translate this query into Propositional Logic as “ $(\neg B)$ ” and form the negated goal as “(B)”.

Your resulting knowledge base (KB) plus the negated goal (in CNF clausal form) is:

$(A B) \quad (A C) \quad (A D) \quad (B C) \quad (C D)$
 $((\neg A) (\neg B)) \quad ((\neg B) (\neg C)) \quad ((\neg C) (\neg D))$
 (B)

Write a resolution proof that Betty had a blue fish.

For each step of the proof, fill in the first two blanks with CNF sentences from KB that will resolve to produce the CNF result that you write in the third (resolvent) blank. The resolvent is the result of resolving the first two sentences. Add the resolvent to KB, and repeat. Use as many steps as necessary, ending with the empty clause. The empty clause indicates a contradiction, and therefore that KB entails the original goal sentence.

The shortest proof that I know of is only four lines long. (A Bonus Point is offered for a shorter proof.) Longer proofs are OK provided they are correct. Think about it, then find a proof that mirrors how you think. Obviously, the four of them must have had two red fish and two blue fish among them. Obviously, if Amy, Cindy, and Diane had two red fish among them, then Betty must have had a blue fish.

Resolve _____ with _____ to produce: _____

Resolve _____ with _____ to produce: _____

Resolve _____ with _____ to produce: _____

Resolve _____ with _____ to produce: _____

Resolve _____ with _____ to produce: _____

Resolve _____ with _____ to produce: _____

Resolve _____ with _____ to produce: _____

Resolve _____ with _____ to produce: _____

Resolve _____ with _____ to produce: _____

Resolve _____ with _____ to produce: _____

6. (15 pts total, 3 pts each) English to FOL Conversion. For each English sentence below, write the FOL sentence that best expresses its intended meaning. Use $\text{Person}(x)$ for “x is a person,” $\text{Game}(x)$ for “x is a game,” and $\text{Plays}(x, y)$ for “x plays y.”

The first one is done for you as an example.

6.a. “Every person plays every game.”

$$\forall x \forall y [\text{Person}(x) \wedge \text{Game}(y)] \Rightarrow \text{Plays}(x, y)$$

6.b. (3 pts) “Some person plays some game.”

6.c. (3 pts) “For every person, there is a game that the person plays.”

6.d. (3 pts) “For every game, there is a person who plays that game.”

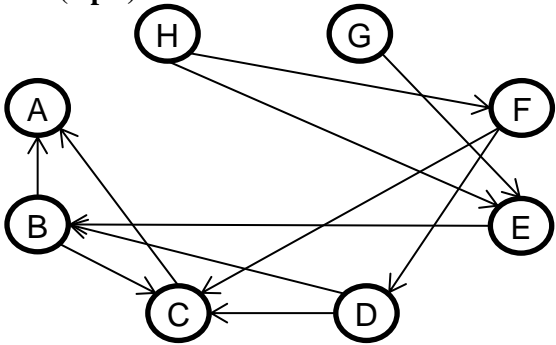
6.e. (3 pts) “There is a game that every person plays.”

6.f. (3 pts) “There is a person who plays every game.”

****** TURN PAGE OVER AND CONTINUE ON THE OTHER SIDE ******

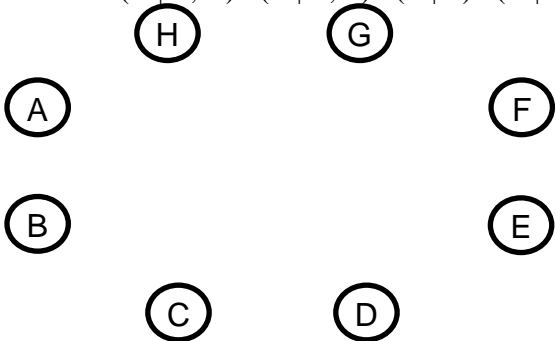
7. (12 pts total, 4 pts each) BAYESIAN NETWORKS.

7.a. (4 pts) Write down the factored conditional probability expression corresponding to this Bayesian Network:

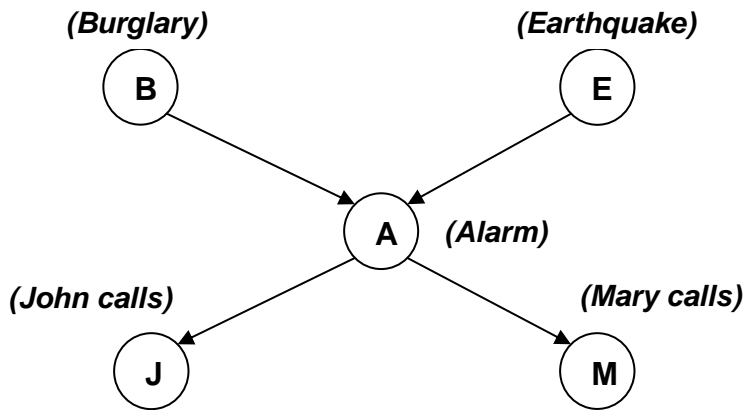


7.b. (4 pts) Draw the Bayesian Network corresponding to this factored conditional probability expression:

$$P(A | C, D) P(B | C, E) P(C | E) P(D | E, F, G) P(E | H) P(F | G, H) P(G) P(H | G)$$



7.c. (4 pts) Shown below is the Bayesian network corresponding to the Burglar Alarm problem, i.e., $P(J, M, A, B, E) = P(J | A) P(M | A) P(A | B, E) P(B) P(E)$. This is Fig. 14.2 in your R&N textbook.



P(B)
.001

P(E)
.002

B	E	P(A)
t	t	.95
t	f	.94
f	t	.29
f	f	.001

A	P(J)
t	.90
f	.05

A	P(M)
t	.70
f	.01

Write an expression that will evaluate to $P(J=f \wedge M=t \wedge A=t \wedge B=t \wedge E=f)$. **Express your answer first as the product of symbolic probabilities, then as a series of numbers (numerical probabilities) separated by multiplication symbols.** You do not need to carry out the multiplication. The first probability of each line is done for you as an example.

$$P(J=f \wedge M=t \wedge A=t \wedge B=t \wedge E=f)$$

$$= P(J=f | A=t) * \underline{\hspace{10em}}$$

$$= .10 * \underline{\hspace{10em}}$$

8. (10 pts total, -1 pt each error, but not negative) Search Properties. Fill in the values of the four evaluation criteria for each search strategy shown. Assume a Tree Search where b is the finite branching factor; d is the depth to the shallowest goal node; m is the maximum depth of the search tree and may be infinite; l is the depth limit; step costs are identical and equal to some positive ϵ ; in bidirectional search both directions use breadth-first search.

Note: These assumptions are the same as in Figure 3.21 of your textbook.

	Complete?	Time complexity	Space complexity	Optimal?
Depth-First				
Breadth-First				
Uniform-Cost				
Depth-Limited				
Iterative Deepening				
Bidirectional (if applicable)				

****** THIS IS THE END OF THE FINAL EXAM ******