

Final Examination

CS 121: Introduction to Artificial Intelligence
Stanford University
Summer 2006

Examination Instructions

You have 3 hours to complete this exam. The complete exam is worth 35% of your total course grade. There are 9 sections, worth a total of 35 points. For each question we have specified the number of points that it is worth. Since you have 180 minutes for the complete exam, you should spend approximately 20 minutes on each problem; please pace yourself and watch the clock.

Please show your work in the space provided, because we will award partial credit. If you need more space, you may use additional sheets of paper, provided you turn them in with your exam.

Stanford University Honor Code

The Honor Code is the University's statement on academic integrity written by students in 1921. It articulates University expectations of students and faculty in establishing and maintaining the highest standards in academic work:

- The Honor Code is an undertaking of the students, individually and collectively:
 - that they will not give or receive aid in examinations; that they will not give or receive unpermitted aid in class work, in the preparation of reports, or in any other work that is to be used by the instructor as the basis of grading;
 - that they will do their share and take an active part in seeing to it that others as well as themselves uphold the spirit and letter of the Honor Code.
- The faculty on its part manifests its confidence in the honor of its students by refraining from proctoring examinations and from taking unusual and unreasonable precautions to prevent the forms of dishonesty mentioned above. The faculty will also avoid, as far as practicable, academic procedures that create temptations to violate the Honor Code.
- While the faculty alone has the right and obligation to set academic requirements, the students and faculty will work together to establish optimal conditions for honorable academic work.

Signature

I attest that I have not given or received aid in this examination, and that I have done my share and taken an active part in seeing to it that others as well as myself uphold the spirit and letter of the Stanford University Honor Code.

Name (printed): _____ SUNet ID: _____

Signature: _____

1. **True/False Questions.** For each statement circle either TRUE or FALSE. Half credit will be given for statements that are left completely blank, and thus it is **not in your best interest to guess randomly**. (0.25 point each, for a total of 4 points)


1	The goal of AI is to design systems that can act like humans.	TRUE	FALSE
2	The Turing test evaluates a system's ability to act rationally.	TRUE	FALSE
3	A search problem must name exactly one state as being the goal.	TRUE	FALSE
4	A finite problem graph can give rise to an infinite search tree with DFS.	TRUE	FALSE
5	Iterative Deepening Search is a way to make DFS optimal for search problems with a constant step cost.	TRUE	FALSE
6	Greedy hill climbing search is often yields better solutions than simulated annealing since it is not stochastic.	TRUE	FALSE
7	The arc-consistency algorithm is only useful if it is run after every variable assignment in CSP search.	TRUE	FALSE
8	The minimum remaining values heuristic orders the values to consider in CSP search.	TRUE	FALSE
9	Given a set of variables, the Bayesian network with no arcs corresponds to one particular probability distribution over those variables.	TRUE	FALSE
10	Every variable in a Bayes net is independent of all of its descendents given its children.	TRUE	FALSE
11	Each probability distribution over a set of variables corresponds to a particular Bayesian network over those variables.	TRUE	FALSE
12	A Naive Bayes model requires $O(n)$ parameters, where there are n features.	TRUE	FALSE
13	The Naive Bayes assumption is that that the observations are independent of the class variable.	TRUE	FALSE
14	In an HMM, the current observation is independent of previous observations given the current state.	TRUE	FALSE
15	In the HMM filtering task, we look for the most likely state sequence given a sequence of observations.	TRUE	FALSE
16	The Bellman equations completely define the value function.	TRUE	FALSE
17	An MDP can be seen as an HMM with the addition of reward and action selection nodes.	TRUE	FALSE
18	If a sentence α is not satisfiable, then the sentence $\neg\alpha$ is valid.	TRUE	FALSE
19	If a sentence α is not valid, then the sentence $\neg\alpha$ is satisfiable.	TRUE	FALSE
20	The following inference procedure is sound and complete for FOL: propositionalize the knowledge base and perform truth table inference.	TRUE	FALSE

2. **Foundations.** (2 points)

(a) Briefly describe the Performance, Environment, Actions, and Sensors (PEAS) for an autonomous securities (e.g., stocks, bonds, etc.) trading agent.

(b) What's wrong with the term "artificial intelligence" to describe this field?

3. **Search.** (2 points)




(a) Consider the problem of searching from state a to state i in the directed graph above. Begin by tracing the DFS algorithm, and draw the search tree that results. Assume that the successor function only works in the direction of the directed arcs. Furthermore, assume that successor states are added to the queue in alphabetical order (a is added before b). To make the tree clearer, circle nodes when they are pulled off the queue.

(b) Now trace the behavior of the BFS algorithm, and draw the search tree that results, in the same manner as above.

- (c) Now suppose we assume that all path costs are 1, and that we are given the heuristic in the following table. Run the A* search algorithm.


State	h
a	3
b	1
c	1
d	1
e	2
f	2
g	2
h	2
i	0

- (d) Is the heuristic admissible?



- (e) Consider the *Towers of Hanoi* puzzle shown above. It is an ancient puzzle, consisting of a collection of rings of different sizes, and three posts mounted on a base. At the beginning all the rings are on the left-most post as shown above, and the goal is to move them all to the rightmost post, by moving one ring at a time from one post to another. But, at no time may a larger ring be placed on top of a smaller one. Formalize the problem of getting from the beginning state to the goal state as a search problem.

4. **Constraint Satisfaction Problems.** (2 points)



- (a) Recall the map coloring problem from class. Suppose we'd like to color the map of Western Europe shown above using only 3 colors: red, green, and blue. Formalize this as a CSP.

- (b) Suppose France must be colored red. Trace the steps of a single run of the *arc-consistency* algorithm to remove inconsistent values from the domains of all the variables. Show your work below.

5. Probability and Bayesian Networks. (2 points)

X	Y	Z	$P(x, y, z)$
true	true	true	0.018
true	true	false	0.002
true	false	true	0.072
true	false	false	0.008
false	true	true	0.252
false	true	false	0.378
false	false	true	0.108
false	false	false	0.162

(a) Consider the full joint probability distribution over the variables X , Y , and Z shown above. What's the marginal probability distribution $P(X)$?

(b) Compute the conditional probability tables $P(Y|X)$ and $P(Z|X)$.

(c) Are Z and Y conditionally independent given X ?

(d) Given your answers to the questions above, design a Bayesian network to represent this joint probability distribution more compactly.

(e) How many independent parameters does your network have?

6. **Machine Learning.** (2 points)

- (a) Suppose we'd like to train a Naive Bayes classifier for a binary classification task. Suppose we have a binary class variable *Happy*, and that there are three binary features, *Sunny*, *Weekend*, and *Productive*. Estimate the parameters of the Naive Bayes model using maximum likelihood estimation from the following training set:

Sunny	Weekend	Productive	Happy
true	true	true	false
true	true	false	true
true	false	false	true
false	true	true	false
false	false	true	true
false	false	false	false

- (b) Compute the probability distribution specified by your learned Naive Bayes model for the *Happy* variable for the two test examples below. What is the most likely value of the *Happy* variable?

Sunny	Weekend	Productive	Happy
false	true	false	?
true	false	true	?

(c) Now train the perceptron on the same training set to predict the *Happy* variable, where we map *true* to the feature value 1 and *false* to the feature value 0. Don't forget to add a bias feature! Do one epoch of training, with $\alpha = 1$. Show your work, and the final weight vector.

(d) Compute the scores assigned to the two possible values of the variable *Happy* for the examples in the test set. What value does it select?

7. **Hidden Markov Models.** (2 points)

- (a) Suppose we are developing software to track airplanes using blips observed on a radar screen. In particular, radar is a *noisy* sensing device, which reports radio waves that bounce off of objects in the sky. When objects are present, they usually (but not always) reflect the waves, and a blip shows up on the screen in the approximate location. The radar sensor rotates around a central axis to pick up readings in every direction, and it makes a full rotation every 10 seconds. Thus the blips on the screen are updated every 10 seconds. Assuming we can discretize the sensible space around us into a 1000×1000 matrix, formalize the relationship between the airplane locations and the blips on the radar screen as an HMM.

- (b) What algorithm would you use to maintain the best possible estimates of the current airplane positions, given the noisy evidence?

- (c) Suppose we are given the following HMM designed to extract animal names from tokenized text, where each hidden state X_t can take on the values *background* or *animal*, and where the evidence variables E_t can take on the value of any token (word or punctuation) in the vocabulary. Here are the HMM parameters:

Starting Distribution

$P(X_0 = \text{background})$	$P(X_0 = \text{animal})$
0.9	0.1

Transition Model

X_{t-1}	$P(X_t = \text{background} X_{t-1})$	$P(X_t = \text{animal} X_{t-1})$
background	0.9	0.1
animal	0.5	0.5

Emission Model

X_t	$P(E_t)$					
	a	chuck	could	how	much	wood
background	0.2	0.1	0.2	0.2	0.2	0.1
animal	0.1	0.3	0.1	0.1	0.1	0.3

- (d) Compute the most likely sequence of states corresponding to the following input sequence. Show the trellis diagram you used in the computation. Did the HMM successfully label the words “wood chuck” as an animal name?

how much wood could a wood chuck chuck

8. **Markov Decision Processes and Reinforcement Learning.** (2 points)

- (a) Suppose you are given a simple 3 state MDP, specified with the following transition function T and reward function R . Run the value iteration algorithm until the policy converges in the table provided below. Use the discount factor $\gamma = 0.5$. When choosing the maximum valued action for each state, break ties by choosing left before right.

Reward Function

S	$R(S)$
a	2
b	4
c	-10

Transition Model

S_{t-1}	A_{t-1}	$P(S_t = a S_{t-1}, A_{t-1})$	$P(S_t = b S_{t-1}, A_{t-1})$	$P(S_t = c S_{t-1}, A_{t-1})$
a	left	0	0.2	0.8
a	right	0	0.8	0.2
b	left	0.8	0	0.2
b	right	0.2	0	0.8
c	left	0.2	0.8	0
c	right	0.8	0.2	0

Value Iteration Table

Iteration	$U(a)$	$U(b)$	$U(c)$	$\pi(a)$	$\pi(b)$	$\pi(c)$
1	0	0	0	left	left	left
2						
3						
4						
5						
6						
7						

- (b) Remember the demo we saw in class of a simulated robot learning to locomote along a line with a single arm (that could rotate around an axis, and bend in one place). Formalize the robot's environment/task as an MDP. You can assume that we can discretize the angular positions of the two arm joints. What is the state space? What is the reward function? What are the possible actions, and what is the transition function?

9. **Logic.** (2 points)

- (a) Use a truth table to prove whether the following propositional logic sentence is valid, satisfiable, or unsatisfiable:

$$(p \wedge (p \rightarrow q) \wedge (q \rightarrow r)) \leftrightarrow (q \wedge r)$$

- (b) Write down first order logic representations of the following English sentences. Use this set of predicates: *IsStudentOf*, *IsSmart*, *IsAdvisor*, *GetsJob*.

Anna is a student of Bill.

All students of a smart advisor are smart.

Bill is smart.

Smart students get jobs.

(c) Now propositionalize the first order logic sentences you came up with above. How many propositions result?

(d) Suppose we wanted to find out whether the following propositionalized sentence is entailed by the other propositionalized sentences above. If we used truth table inference, how many rows would the truth table have?

GetsJob(Anna)

(e) Use the propositional forward chaining algorithm instead, with the modus ponens rule, to show that the sentence is entailed.