Name: ______________________________

Final

COS302, Fall 2001, Michael L. Littman

There are six questions. Choose 4 to do for 25 points each by circling their problem numbers. This gives a total of 100. All uncircled problems are extra credit and are worth 1/5 value. Thus, maximum score with extra credit is 110.

No notes.
If more than four problem numbers are circled, I'll uncircle some at random.
Keep your answers as short as possible.
Write and sign the honor code below.
Please read all instructions carefully!
Find me in the back if you need assistance.
Partial credit will be awarded if you show your work.
If you feel yourself getting stuck on a problem, move on and return to it later.

You can do it!

1. **Competitive Binary Search**

Two players are each given a random number from one to ten. Players alternate asking questions to try to determine the other player's number. A player may ask a rule-out question "Is your number less than x?" or a final question "Is your number x?". The game ends after a final question with the winner being the asker if the answer was "yes" or the askee if the answer was "no".

In analyzing this game, note that all that matters is the number of choices that have not been ruled out for each of the players. The state of the game is a pair of numbers (i, j) corresponding to the number of choices left for the first and second player. The "rule out" actions are the split size: any number k from 1 to $i-1$, splitting the i numbers into a set of size k and a set of size $i-k$.

To solve this, we need to know for each state, the expected minimax value of the game. The table below gives this value for i's turn for all states.

i	j: 1	2	3	4	5	6	7	8	9	10
1	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
2	0.500	0.500	0.667	0.750	0.800	0.833	0.857	0.875	0.889	0.900
3	0.333	0.500	0.556	0.583	0.667	0.722	0.762	0.792	0.815	0.833
4	0.250	0.500	0.500	0.562	0.600	0.667	0.714	0.750	0.778	0.800
5	0.200	0.400	0.467	0.500	0.560	0.600	0.629	0.650	0.689	0.720
6	0.167	0.333	0.444	0.500	0.533	0.556	0.595	0.625	0.648	0.667
7	0.143	0.286	0.429	0.464	0.514	0.524	0.551	0.571	0.603	0.629
8	0.125	0.250	0.417	0.438	0.500	0.500	0.536	0.562	0.583	0.600
9	0.111	0.222	0.370	0.417	0.467	0.481	0.508	0.528	0.556	0.578
10	0.100	0.200	0.333	0.400	0.440	0.467	0.500	0.525	0.544	0.560

(A). (10 points) It's i's turn and i has 10 numbers left while j has only 9. Compute the expected minimax value for i choosing a 6–4 split and a balanced 5–5 split. Circle the best action.

(B). (15 points) Write the equation used to compute the $[i, j]$ entry in the table. Don't forget the base case!

2. **Dreidel Learning**

"I have a little dreidel, I made it out of clay. I made it sort of lumpy, it rarely lands on hey."

I spun my dreidel (random number generator) 10 times to create a training set, then 10 more times to create a validation set. Here are the outcomes:

val	name	train	validate
1	gimmel	1	0
2	hey	0	2
3	nun	7	5
4	shin	2	3

(A). (5 points) What's the maximum likelihood model based on the training data?

(B). (5 points) How likely is the validation data using the answer from Part A?

(C). (5 points) We can use Laplace smoothing to make the validation data more likely and learn a better model. Give the smoothed model.

(D). (10 points) How likely is the validation data using the answer from Part C? The probability is better than that of the unsmoothed model (Part B) and the "uninformed" model (uniform probabilities).

3. **Monte Hall's Problem**

It's the Price is Right! There are three doors, one hides a prize, and Monte knows which one. The contestant chooses a door. Monte now has a choice. He can either do nothing, or show the contestant which of the other two doors does *not* have the prize. The contestant may now do nothing or switch to one of the other doors. Now, the doors are all opened. If the contestant picked the door with the prize, she wins \$2. If the contestant switched away from the door with the prize, she loses \$1. Otherwise, she gets nothing.

Monte's goal is to minimize the contestant's winnings.

(A). (15 points) Draw the game matrix. Rows should be Monte's strategies, columns should be the contestant's strategies, cells should contain the contestants' expected winnings.

(B). (10 points) What strategy (possibly probabilistic) should Monte adopt to minimize the contestant's winnings? What are the resulting expected winnings?

4. **Part of Speech Tagging**

Here's a bigram part of speech model.

$\Pr(w|t)$:

DET	the (1.00)		
NOUN	man (.53)	saw (.02)	her (.45)
POS	her (1.00)		
VERB	saw (.01)	send (.87)	man (.12)
EOS	. (1.00)		

$\Pr(t'|t)$:

	DET	NOUN	POS	VERB	EOS
DET	.00	.95	.00	.01	.04
NOUN	.01	.17	.01	.04	.77
POS	.00	.98	.00	.01	.01
VERB	.17	.08	.51	.03	.21

(A). (10 points) Find the most likely sequence of tags for the sentence below. Use the table to fill in Viterbi probabilities. Use four-digit accuracy.

tags:

	the	man	saw	her	.
DET					
NOUN					
POS					
VERB					
EOS					

(B). (15 points) Same question, different sentence.

tags:

	the	man	saw	her	saw	.
DET						
NOUN						
POS						
VERB						
EOS						

5. **Regional Dialects**

(A). (10 points) Draw a decision tree on the right that correctly classifies the points in this diagram. Each node can contain an inequality expression for a single variable.

y \ x	.0	.1	.2	.3	.4	.5	.6	.7	.8	.9
.0						$-$				
.1										$-$
.2		$+$								
.3							$+$			
.4	$-$									
.5										
.6			$+$							
.7										
.8								$+$		
.9										

(B). (5 points) Can a perceptron correctly classify the points above? Explain.

(B). (10 points) Classify all the points in the diagram according to the decision tree given (left = true, right = false).

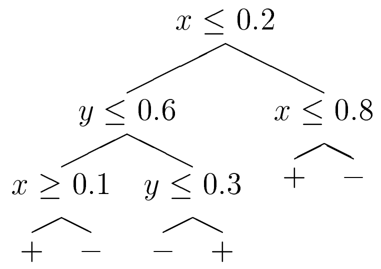

y \ x	.0	.1	.2	.3	.4	.5	.6	.7	.8	.9
.0										
.1										
.2										
.3										
.4										
.5										
.6										
.7										
.8										
.9										

6. **Indoor/outdoor**

We'd like to classify photographs as being either "indoor" or "outdoor" on the basis of the words in the captions. We decide to use a unigram model and a Bayesian classifier, which we'll train using EM.

Here's the initial probability model.

Pr(inside) = 0.4

Pr(outside) = 0.6

	x	y: snow	sport	chair	speech
Pr(y|x):	inside	.1	.2	.3	.4
	outside	.3	.5	.1	.1

(A). (15 points) For each of the two-word captions, compute the probability that each is an inside picture or an outside picture based on the model above. (E-step)

	Pr(inside)	Pr(outside)
snow sport		
sport speech		
chair chair		

(B). (10 points) From the probability you computed, derive a new maximum liklihood model. (M-step)

Pr(inside) =

Pr(outside) =

	x	y: snow	sport	chair	speech
Pr(y|x):	inside				
	outside				