4511W, Spring-2020
ASSIGNMENT 4:
Assigned: 03/24/20 Due: ~~03/30/20~~ 04/01/20 at 11:55 PM (submit via Canvas, you may take a picture of handwritten solutions, but you must put them in a pdf) Submit only pdf or txt files

Written/drawn:

Problem 1. (20 points)
Assume you have a binary tree. If you ran Monte-Carlo tree search for 100 rollouts (random playouts), what is the maximum times you could choose the left branch from the root (i.e. go down left from the initial node)? (You should have some math to backup your answer, not just a ballpark answer.) (Clarification: using UCB values as discussed in class, rather than some other way)

Problem 2. (15 points)
If you do not know how to play checkers, go read this: https://simple.wikipedia.org/wiki/Checkers .
Assume you have a a 4x6 board (with players starting on the length 4 sides) as such:

X				O	
	X				O
X				O	
	X				O

Distance from X side:

0	1	2	3	4	5

Assume you have a mid-state evaluation as:

$$\left(\sum_{\forall p \in \{\text{all your pieces}\}} DistanceFromYourSideTo(p) \right) - \sum_{\forall q \in \{\text{opponent's pieces}\}} 1$$

(1) What is the value of the following state for the X player:

X				X	
			O		O
X		X			
					O

(2) Would it work to use the mid-state evaluations in conjunction with minimax to find the best move? Explain your reasoning clearly.

Problem 3. (20 points)
Assume you have the following payoff matrix (first number for the row value (up-down), second for the column (left-right)):

6, 5	3, 5	3, 8
5, 3	4, 6	3, 3
1, 7	2, 3	5, 4

(1) What are the pure-strategy Nash-equilibrium? Find at least one mixed-strategy Nash-equilibrium

(2) What are the pure-strategy Pareto-optimum? What are the mixed-strategy Pareto-optimum (think carefully about this)?

Problem 4. (15 points)
Assume you and an opponent are playing the following Prisoner's dilemma game:

	P2 confess	P2 lie
P1 confess	P1 = -8, P2 = -8	P1 = 0, P2 = -10
P1 lie	P1 = -10, P2 = 0	P1 = -1, P2 = -1

(1) You are going to play this game for 3 rounds with/against a tit-for-tat agent. The “cooperative” action in this game is to “lie”, while the “competitive” action is to “confess”. The value of a terminal state is the sum of the rewards from the 3 games along the path. What is the best way to play?

(2) You do not need to actually search the entire tree to answer the first part of this question. Which branches could you have skipped? What algorithm are you running in disguise?

Problem 5. (15 points)
Suppose you have a CSP with three variables: a, b, and c. These variables all start with single digit domains (i.e. 0 to 9). Assume the constraints are:

- $b \cdot c < a^2$
- $2a > b + 4$
- $b + c = 10$
- $a < c$

(1) Find the domains for a, b, and c that are 2-consistent.

(2) Which pairs from part (1) are 3-consistent? What are all the possible solutions for a, b and c? (You can find the answer to these in whichever order you find easiest.)

Programming:

Problem 6. (15 points)
The “Zebra” problem is a classic constraint satisfaction problem defined as a set of 5: houses, people (of nationality), pets, cigars, drinks and colors. House 1 is on the far left with house 5 on the far right (house 3 is in the middle). Each house (numbered) has a single person (nationality), pet, cigar, drink and color associated with the house. No two houses share the same of any property (for example, each house has its own unique pet). The rules are:

- There are five houses.
- The Englishman lives in the red house.

- The Spaniard owns the dog.
- Coffee is drunk in the green house.
- The Ukrainian drinks tea.
- The green house is immediately to the right of the ivory house.
- The Old Gold smoker owns snails.
- Kools are smoked in the yellow house.
- Milk is drunk in the middle house.
- The Norwegian lives in the first house.
- The man who smokes Chesterfields lives in the house next to the man with the fox.
- Kools are smoked in the house next to the house where the horse is kept.
- The Lucky Strike smoker drinks orange juice.
- The Japanese smokes Parliaments.
- The Norwegian lives next to the blue house.

The question is then: Who owns the Zebra? And who drinks water?

(5/15 points) This classic problem is what is already put in as the Zebra problem. Run the backtracking_search() on this problem and report the answer to the two questions above. (It is fine to use the default parameters for backtracking_search().) (Note: these are in csp.py and tests/test_csp.py)

(10/15 points) Modify the problem to match the Zebra problem below (with names, sports, transportation and lawns) and report: Who has flowers in their yard? Who likes to watch Starcraft2?

- There are five houses.
- Albert likes baseball.
- Dietfried has a lawn with cleanly cut grass.
- Virgilijus enjoys watching rugby.
- Gallchobhar gets to work by walking.
- Bricius lives next to Virgilijus.
- The owner of the house with an Astroturf lawn likes baseball.
- The house with flowers in the lawn is between the house with trees and the house with rocks in the lawns.
- The person who likes baseball goes to work on a motorcycle.
- The 2nd house owners like to watch debates.
- The owner of the 4th house goes to work on a bicycle.
- The person who likes soccer takes a bus to work.
- The house with trees on the lawn to the right of the house with an owner who drives a car to work.
- Gallchobhar lives to the right of the person who likes Rugby.