15-381 Spring 06 Assignment 3: Solution

Questions to Rong Yan(yanrong@cs.cmu.edu)

Out: 2/21/06 Due: 3/7/06

Name: ______________________ Andrew ID: ______________________

Please turn in your answers on this assignment (extra copies can be obtained from the class web page). This assignment must be turned in at the beginning of class at 1:30pm on March 7. Please write your name and Andrew ID in the space provided on the first page, and write your Andrew ID in the space provided on each subsequent page. This is worth 5 points: if you do not write your name/Andrew ID in every space provided, you will lose 5 points.

This assignment has no coding problems

Late policy. Your assignment is due at 1:30pm on 3/7. Submitting your work late will affect its score as follows:

- If you submit it after 1:30pm on 3/7 but before 1:30pm on 3/8, it will receive 90% of its score.
- If you submit it after 1:30pm on 3/8 but before 1:30pm on 3/9, it will receive 50% of its score.
- If you submit it after 1:30pm on 3/9, it will receive no score.

Collaboration policy. You are to complete this assignment individually. However, you are encouraged to discuss the general algorithms and ideas in the class in order to help each other answer homework questions. You are also welcome to give each other examples that are not on the assignment in order to demonstrate how to solve problems. But we require you to:

- not explicitly tell each other the answers
- not to copy answers
- not to allow your answers to be copied

In those cases where you work with one or more other people on the general discussion of the assignment and surrounding topics, we ask that you specifically record on the assignment the names of the people you were in discussion with (or "none" if you did not talk with anyone else). This is worth five points: for each problem, space has been provided for you to either write people's names or "none". If you leave any of these spaces blank, you will lose five points. This will help resolve the situation where a mistake in general discussion led to a replicated weird error among multiple solutions. This policy has been established in order to be fair to the rest of the students in the class. We will have a grading policy of watching for cheating and we will follow up if it is detected.

1 The Robot Motion Planning Problem (15 points)

References (names of people I talked with regarding this problem or "none"):

____

Consider the robot planning problem illustrated in the following figure.

1.1 (5 points)

Draw (approximately) the path found by using potential fields.

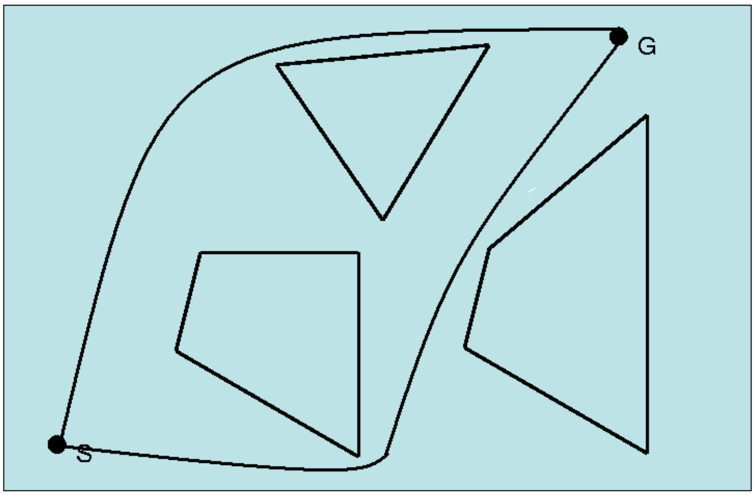

Solution: *Since we do not define the exact form of the potential field function, we accept multiple correct answers for this problem. Above are two possible paths.*

1.2 (5 points)

Draw (approximately) the path found by using approximate cell decomposition (there are several possibilities). Briefly discuss what are the advantages and disadvantages of the approximate cell decomposition technique.

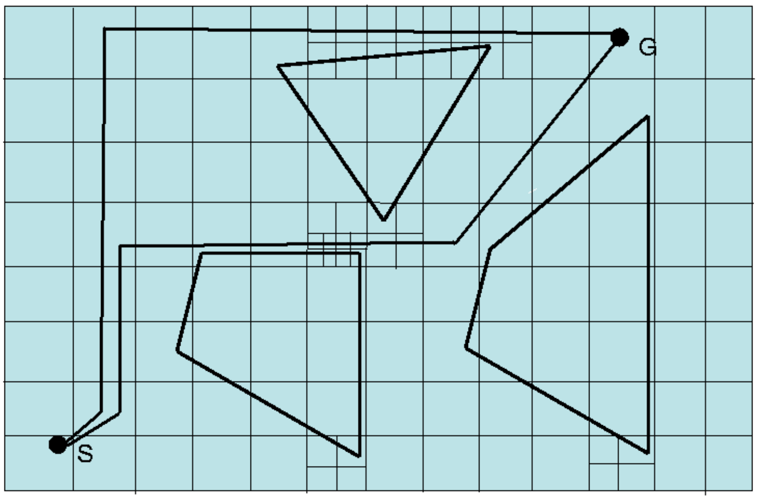

Solution:

Advantage: 1. limited assumption on obstacle; 2. used in practice; 3. find obvious solutions quickly.

Disadvantage: 1. No clear notion of optimality; 2. tradeoff of completeness/computation; 3. difficult to be used in high dimension.

1.3 (5 points)

Draw (exactly) the path found by using the visibility graph technique. Draw the initial visibility graph. Briefly discuss what are the advantages and disadvantages of the visibility graph technique.

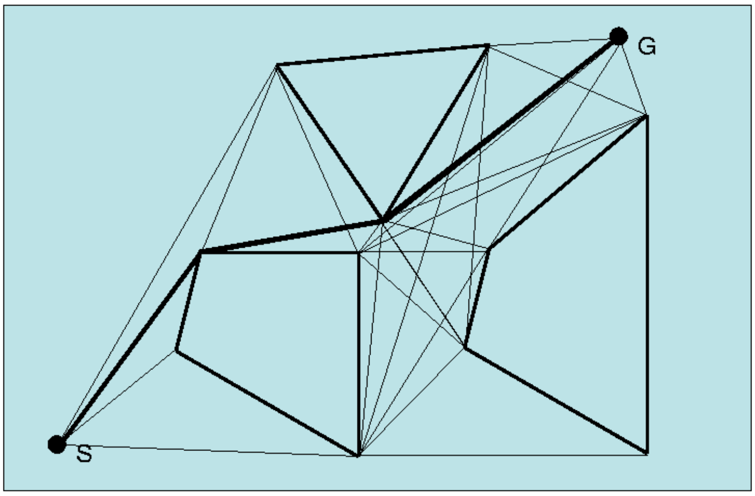

Solution:

Advantage: 1. always find the shortest path; 2. easy to implement.

Disadvantage: 1. try to stay close to the obstacle; 2. easily lead to collision; 3. complicate in high dimension.

2 The Game Play Problem(40 points)

References (names of people I talked with regarding this problem or “none”):

____

Gomoku, go-moku, or gobang is an abstract strategy board game with two players. It is traditionally played with black and white stones on a 19 * 19 go board. One player has black stones and the other has white stones. To simplify, we assume the game are played on a 10*10 board with crosses(x) representing black stones and circles(o) representing white stones. In this game, black plays first and players alternate in placing a stone of their color on an empty space. The winner is the first player to get an unbroken row of five stones horizontally, vertically, or diagonally. Some examples of the winning situations for black stones are shown as follows. Please see http://en.wikipedia.org/wiki/Gomoku for more details about this game. This game is an example of 2-player zero-sum deterministic game of perfect information.

		x						x	
		o	x					x	
			o	x	o			x	
				o	x			x	
						x		x	
			o	o			o		
		x	x	x	x	x			
	o	o			o				

2.1 (5 points)

Any two-player zero-sum deterministic game can be represented by the following quintuple: (S, I, Succ, T, V), where S is the entire space of game states, I is the initial state, Succ is the successor function, T are the terminal states, V maps from terminal states to its payoff/utility. Please describe what is each element of the quintuple in the Gomoku game.

Solution:

S: all possible stone configurations on the board, where the difference between the numbers of black stones and white stones is no more than 1.

I: an empty board.

Succ: place an additional stone on the previous board state.

T: either player has five stones connected horizontally, vertically or diagonally.

V: maps black(white) wins to 1, white(black) wins to -1, a draw to 0.

2.2 (10 points)

Design two different heuristic evaluation functions for this game given the current state of the board. In this problem, the evaluation function should be able to handle the intermediate states of the game. For example, you cannot define "+1 for black wins, -1 for white wins, 0 otherwise" because it cannot provide useful information for the intermediate states. (Hints: you can consider the number of connected stones on the board.)

Solution: *There are many ways to define the heuristic evaluation functions. The following functions are three examples of evaluation functions for the player using black stone.*

1. $100 * SB_5 + 6 * SB_4 + 2 * SB_3$

2. $100 * SB_5 + 6 * SB_4 + 2 * SB_3 - 100 * SW_5 - 6 * SW_4 - 2 * SW_3$

3. $100 * PB_4 + 6 * PB_3 - 100 * PW_4 - 6 * PW_3$

where $SB_i(SW_i)$ is the number of i-connected black(white) stones on the board, $PB_i(PW_i)$ is the number of "live" i-connected black(white) stones. "Live" connected stones means no other stones are placed on both end of the connected stones. Therefore, a live 4-connected stone will guarantee to win.

In the following discussions, let us assume we have arrived to the following state of board. Player x is going to move next.

					o				
				o		x			
			o			x			
		o		x	x	x		o	
	x					o			

2.3 (10 points)

Solve the whole game and find the next move for x with the mini-max algorithm. Rather than using heuristic evaluation functions, we will evaluate the board state based on the simple win/loss between x and o (i.e., +1 for x wins, -1 for x loses and 0 otherwise). For the purpose of grading, you only need to consider the nodes labeled in the following figure (from 1 to 4) to expand and must use the ascending order for expansion, i.e., on each level of expansion, you must expand the lowest remaining numbered node before expanding the higher numbered nodes. Any node where a player have 5 stones connected is the leave node and has no children.

You do *not* need to draw the entire game tree in this problem, but you need to draw the best path found by the mini-max algorithm (if there are several best paths, you should show the first one expanded). For each node in the path, indicate the index of the node expanded and the evaluated value computed/propagated by the algorithm. For each node, also indicate if it is from a max level or a min level. Moreover, please answer: what should be x's first move? how many nodes does the minimax algorithm need to expand in this case?

						1			
					o	3			
				o		x			
			o			x			
		o	2	x	x	x	4	o	
	x					o			

Solution: *The best path is (move 1, MAX level, value 0) → (move 3, MIN level, value 0) → (move 2, MAX level, value 0) → (move 4, MAX level, value 0). The first move for x should be 1.*

There are some ambiguities in the last question. If you interpret "expand" as creating children nodes, the number should be 25 or 26 (depends whether you count the initial node or not). If you interpret "expand" as visiting a node, the number should be 46 or 47. We accept all these answers.

2.4 (10 points)

Solve the whole game with the alpha-beta pruning. The problem settings are the same as the last problem. But you *need* to draw the entire game tree, except for those parts that you can prune (you must leave them out). Indicate which level of the trees are max levels or min levels. For each node in the game tree, indicate the index of the node expanded and the alpha-beta values discovered. In this case, how many nodes does the algorithm need to expand?

Solution: *The game tree is shown as follows. Each cell contains the movement and the update history (from left to right) of alpha/beta, where A means alpha and B means beta. For each node, alpha is initialized to -1 and beta is initialized to 1. Note that we only list alpha for nodes in the MAX levels because beta will not be updated in these nodes. Similarly, we only list beta for the nodes in the MIN levels. Due to the ambiguity of "expand", we accept the answers of 17/18 and 10/11 as the correct answers for the last question.*

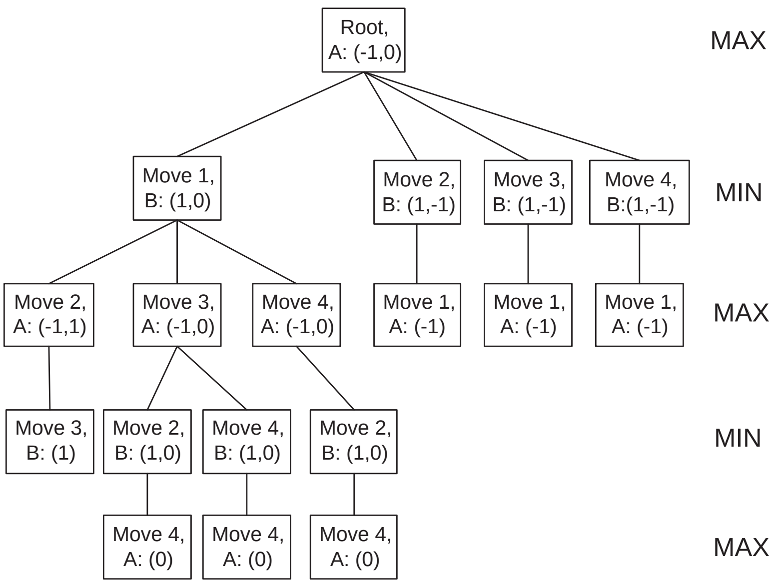

2.5 (5 points)

Now let us consider a different situation. The state of the board is the same as before. But player o need to leave town for his next move (right after the first x's move). So that move will be handled by his representative r. But it is well known that r will just make a random move on the board (from 1 to 4)! Good news is that o will return for his last move. Therefore, the entire moving sequence is x, r, x, o.

Briefly explain how you can modify the mini-max algorithm accordingly to deal with this change. Under this condition, does player x's first move change? Please explain why.

Solution: *To handle the random move of player o, you should modify the second level of your game tree by computing the expectation of your heuristic functions under a uniform probability distribution for o's movements, instead of computing the minimal value. However, under this condition, player x's first move does not change, because the game value for x to pick movement 1 is still higher than the other movements.*

3 Optimal Mixed Strategy (10 points)

Consider the game with the following matrix: (if you look carefully, this is a generalization of the children's game "scissors, rock, paper")

0	1	-2
-1	0	3
2	-3	0

3.1 (5 points)

Show that the value of that game (i.e., the expected payoff that will be obtained by either player using the optimal mixed strategy) is 0. In fact, this is true for any game for which $m_{ij} = -m_{ji}$, but you don't need to prove the general result.

Solution: *Suppose player 1 plays a mixed strategy that uses strategy i with a probability p_i. Then we should prove the maximin payoff of player 1 is equal to 0, i.e.,*

$$\max_p \min(E_{p1}, E_{p2}, E_{p3}) = 0,$$

where E_{pj} is the expected payoff of player 1, i.e., $m_{1j}p_1 + m_{2j}p_2 + m_{3j}p_3$, given the player 1 uses the mixed strategy p and player 2 uses the pure strategy j. Substituting the numbers into the formula, we have to prove,

$$\max_p \min(-p_2 + 2p_3, p_1 - 2p_3, -2p_1 + 3p_2) = 0.$$

Several ways to prove this proposition:

1. Rigorously, you should first show $\min(-p_2 + 2p_3, p_1 - 3p_3, -2p_1 + 3p_2) \leq 0$. *If this is not true, it means* $E_{p1} = -p_2 + 2p_3 > 0$, $E_{p2} = p_1 - 3p_3 > 0$ *and* $E_{p3} = -2p_1 + 3p_2 > 0$, *which will lead to a contradiction by showing* $0 = 3E_{p1} + 2E_{p2} + E_{p3} > 0$. *Then you can find* $(3/6, 2/6, 1/6)$ *is a triplet to satisfy* $\min(-p_2 + 2p_3, p_1 - 3p_3, -2p_1 + 3p_2) = 0$. *This will complete the proof.*

2. Alternatively, you can consider all the intersection points between the pairs of E_{p1}, E_{p2}, E_{p3}, *i.e., the point of* $(3/6, 2/6, 1/6)$, *and also all the boundary points of the probability distributions, i.e.,* $(1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$. *Prove* $\min(-p_2 + 2p_3, p_1 - 3p_3, -2p_1 + 3p_2) \leq 0$ *on these points and thus it is true for all the probability distributions.*

3. As another way to prove, you can first show that the optimal mixed strategies of player 1 and player 2 are the same given the game is symmetry (strictly speaking, you should prove this is correct, but we will still give you credits if you just hand wave here). Thus, the expected payoff of player 1 in this case is always 0.

3.2 (5 points)

Derive the optimal mixed strategy, which is the same for both players. If you have trouble with question 1, just assume that the value of the game is 0 and proceed with this question.

Solution: (p_1, p_2, p_3)**=(1/2, 1/3, 1/6)**

4 Widget Companies (10 points)

Consider two widget companies, Acme and U.S. Widgets, competing for the same market and each firm must choose a high price ($2 per widget) or a low price ($1 per widget). Here are the rules of the game:

1. At a price of $2, 5000 widgets can be sold for a total revenue of $10000;
2. At a price of $1, 10000 widgets can be sold for a total revenue of $10000;
3. If both companies charge the same price, they split the sales evenly between them;
4. If one company charges a higher price, the company with the lower price sells the whole amount and the company with the higher price sells nothing;
5. Payoffs are profits – revenue minus the $5000 fixed cost.

(Note: This is a very special case of something called the Cournot's game that models price competition. Unfortunately, it is an overly simplified and unrealistic example because we force it to be a zero-sum game.)

4.1 (5 points)

Write down the game matrix.

Solution:

Suppose Acme's strategies correspond to the rows and U.S. Widgets' strategies correspond to the columns. The game matrix for Acme is shown as follows,

	1	2
1	0	5000
2	-5000	0

4.2 (5 points)

Is there a pure strategy solution? Briefly explain.

Solution: *There is one pure strategy solution when both companies are selling the products at the price of $1. This is because both companies will get lower profits if they switch to other strategies in this situtaion.*

5 The Nim Game(10 points)

A very simple version of the game of "Nim" is played as follows. There are 2 players and, at the start, two piles on the table in front of them, each containing 2 matches. In turn, the players take any positive number of matches from one of the piles. The player taking the last match loses.

Now consider a more elaborate version of the game of "Nim" is played as follows. There are 2 players and, at the start, three piles on the table in front of them, each containing 2 matches. In turn, the players take any positive number of matches from one of the piles. The player taking the last match loses. Sketch a game tree. Show that the first player has a sure win.

Solution: *The game tree is sketched as follows. In this tree, we only consider the best strategy for player 1 and do not expand any redundant nodes for player 2. For example, in the second level, we do not expand the nodes of (-, ii, i), since it will have the same outcome as (-, i, ii).*

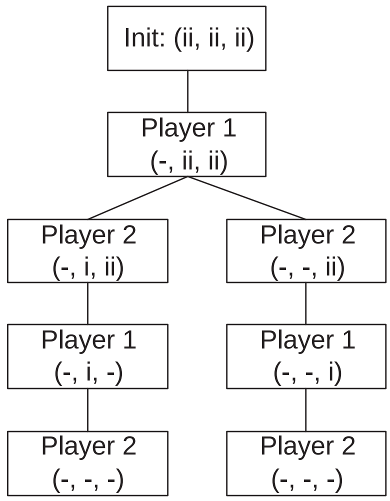

6 The Pin Game(10 points)

A game is played as follows: The two players, i.e., A and B, simultaneously hold up either one or two pins. A wins if there is a match (the number of pins are the same for both players) and B wins otherwise. The amount won is the number of pins held up by the winner. It is paid to the winner by the loser.

6.1 (5 points)

Write down the matrix form of the game and specify the 2 players' possible pure strategies.

Solution: *Suppose A's strategies correspond to the rows and B's strategies correspond to the columns. The game matrix for A is shown as follows,*

	1	2
1	1	-2
2	-1	2

6.2 (2 points)

Explain why there is no pure strategy solution in this game.

Solution: *Because none of these 4 game states are equilibria. For example, if we start with the state (1,1), B will switch to the strategy 2 in order to get a higher payoff, and then A will switch to the strategy 2 for the same reason. This iterative process will not be converged. Therefore, there is not pure strategy solution in this game. BTW, it is also ok to argue that "maximin != minimax" in this game matrix.*

6.3 (3 points)

Compute the mixed strategy solution and the value of the game.

Solution: *Assume A plays a mixed strategy {1:p, 2:(1-p)}, and B plays a mixed strategy {1:q, 2:(1-q)}. The mixed strategy solution is,*

$$p^* = \arg\max_p \min(p - (1-p), -2p + 2(1-p)) = 1/2,$$

$$q^* = \arg\min_q \max(q - 2(1-q), -q + 2(1-q)) = 2/3,$$

and the game value is 0. Note: when you are computing the mixed strategy of q, you need to use the minimax process.

7 Mixed strategy(5 points)

Consider a two-player game in which each player has 4 possible strategies (denoted 1 to 4) with the matrix form shown below. Compute a mixed strategy solution, assuming that each player assigns the same probability to strategies 1, 2, and 3.

	1	2	3	4
1	-1	1	1	-1
2	1	-1	1	-1
3	1	1	-1	-1
4	-1	-1	-1	1

Solution: *Assume player 1 plays a mixed strategy {1:p, 2:p, 3:p, 4:(1-3p)}, and B plays a mixed strategy {1:q, 2:q, 3:q, 4:(1-3q)}. The mixed strategy solution is,*

$$p^* = \arg\max_p \min(p - (1 - 3p), -3p + (1 - 3p)) = 1/5$$

$$q^* = \arg\min_q \max(-q + (1 - 3q), 3q + (1 - 3q)) = 1/5$$