CSCE420: Introduction to Artificial Intelligence
Prior Class Test Questions

The following questions are all questions from prior midterm examinations. They represent more than what is is feasible in 50 minutes under test conditions but, for purposes of study and revision, providing more examples is preferable to fewer.

Important: Your exam assumes that you will have prepared 3 pages of notes. These can be US-letter, double-sided, typeset or written—to be brought with you on Mar 6th. These practice questions can give you a sense of what it might be useful to include in your notes.

Question 1	6
Question 2	5
Question 3	6
Question 4	8
Question 5	15
Question 6	14

Question 1. (6 points)

In the context of intelligent agents there is a notion of a *known* environment. What does this mean, and how would it relate to a web-bot (or web-crawler, web-spider) whose search space is the World Wide Web? (2 pts.)

What is the relationship of quadratic programming to linear programming? Be as precise as you can. (2 pts.)

When is the least-constraining-value heuristic useless in a CSP search? (2 pts.)

Question 2. (5 points)

Construct a map with 5 cities $\{A, B, C, D, E\}$ and 5 roads each of unit length, so that there are two routes from A to E. Give the values of an heuristic function so that depth-first search takes fewer steps to find the optimal path from A to E than A* search. (5 pts.)

Question 3. (6 points)

Assume that alpha-beta pruning reduces the branching factor from b to $\sqrt{b}$. How much deeper (in terms of plies) will an alpha–beta pruned search solve when compared with the standard minimax algorithm, given that both run for the same amount of time. Explain why mathematically. (3 pts.)

You are programming an agent to playing a complex game in which there is only a limited, fixed amount of time to search between moves. Your solution is to run minimax search with alpha–beta pruning until the time is nearly up, and then to make the best move found far. Suppose your agent does this and is about to pick the move. At that time the alpha and beta bounds for the current state have values 3 and 7 respectively. What can the difference between these two values be interpreted? (3 pts.)

Question 4. (8 points)

Suppose you are given a screen-shot of a game of minesweeper in progress. You get a description of cells marked with flags as mines, some numbered with a count of neighboring mines (from zero to eight), and others yet to be opened. You cleverly decide to cast this as a CSP. You use a variable with domain $\{0, 1\}$ for each cell. Cells marked with mines are set to have value 1, those with known numbers are set of have value 0, and you introduce a summation constraint on the 8 neighboring cells. Suppose you use modify your favorite CSP algorithm to count the number of satisfying solutions.

If it tells you that there is only one solution, what do you conclude?

It tells you that there are multiple satisfying solutions. What do you conclude?

Suppose it tells you that there are no satisfying solutions. What do you conclude?

Question 5: Heuristics (15 points)

Consider heuristics where $h(n) = 0$ for all n for which $\text{GOAL}(n) = \text{True}$, i.e., functions that have the value zero for states that are goals. Then the following statements are either true or false

Statement 1: Every admissible heuristic is consistent.
Statement 2: Every consistent heuristic is admissible.

For each statement: (1) identify whether it is true or false. (2) If it is true prove that it is so; otherwise, if it is false, provide a counterexample.

> *Reminder:* Using the notation from the textbook, with n' a successor of state n and a an action, a consistent heuristic is required to have $h(n) \leq c(n, a, n') + h(n')$.

Question 6: Analysis of Games (14 points)

Consider a two-person, zero-sum game of perfect information we call G. Think of a game like checkers, or tic-tac-toe, but unlike those games, in G, more extreme scores are awarded for demolishing the opponent. Thus, the outcomes can be $O = \{-2, -1, 0, 1, 2\}$.

Suppose that agent $\mathcal{A}$ implements alpha-beta pruning correctly; it does a complete search down all the tiers reaching the leaves (i.e., it runs without cut-offs). Additionally, a fellow student has their own algorithm, $\mathcal{S}$, they've built using their own custom insights.

In the following, we report matches from playing G by pairs of players. When we say that the outcome of X *vs.* Y was (o_x, o_y), we mean that agent X played the first move as a maximizing player, that Y played second and was the minimizing player, and, after the game finished, X got outcome $o_x \in O$, and Y got $o_y \in O$. *(Answer each of these questions anew, discarding any inferences drawn from previous answers.)*

6.1 Is it possible to find that $\mathcal{A}$ *vs.* $\mathcal{S}$ yielded the outcome $(-1, 2)$, for the agents respectively? Explain/Interpret. (2 pts)

6.2 Is it possible for $\mathcal{A}$ *vs.* $\mathcal{S}$ to yield $(-1, 1)$, respectively, and for $\mathcal{S}$ *vs.* $\mathcal{A}$ to yield $(1, -1)$?
Explain/Interpret. (2 pts)

6.3 Is it possible for $\mathcal{A}$ *vs.* $\mathcal{S}$ to yield $(1, -1)$, respectively, and for $\mathcal{S}$ *vs.* $\mathcal{A}$ to yield $(-1, 1)$?
Explain/Interpret. (2 pts)

6.4 Is it possible for $\mathcal{A}$ *vs.* $\mathcal{A}$ to yield $(0, 0)$? Explain/Interpret. (2 pts)

6.5 Is it possible for $\mathcal{A}$ *vs.* $\mathcal{A}$ to yield $(2, -2)$ when $\mathcal{A}$ *vs.* $\mathcal{S}$ gave $(1, -1)$? (2 pts)

Both agents now play a third player $\mathcal{T}$.

6.6 Can $\mathcal{A}$ *vs.* $\mathcal{T}$ yield $(1, -1)$, and $\mathcal{T}$ *vs.* $\mathcal{A}$ give $(-1, 1)$, when $\mathcal{S}$ *vs.* $\mathcal{T}$ yielded $(2, -2)$, and also $\mathcal{T}$ *vs.* $\mathcal{S}$ resulted in $(-2, 2)$? Explain/Interpret. (4 pts)