Mathematical Foundations of AI (67686): Take Home Exam

Ariel Procaccia

Deadline: 31.8.08

Instructions: The exam can be done in pairs. However, you may not speak about the exam with anyone besides your partner (if you are doing the exam alone, you may not speak about it with anyone!). The deadline for submission is 31.8.08. Please write your solution in LaTex/Word, and submit the pdf/doc by email. Good luck!

1 Robotic Motion (30 points)

You are given an $n \times n$ grid, where each cell is either blocked or unblocked, and the starting locations of two robots. Initially, each robot can see the 4-neighbors adjacent to its starting location. In each step, *both* robots can move to a 4-neighbor, and the 4-neighbors of the robots' new locations are revealed. This is exactly the setting discussed in Lecture 3 of the course (only with two robots instead of one). The goal is to get our two robots to meet, that is, to move to the same cell, using a minimal number of steps.

1. **(20 points)** For any given constant $\epsilon > 0$, give an $O(n^{1+\epsilon})$-competitive algorithm for the problem.

2. **(10 points)** Give a lower bound of $\Omega(n)$ for the competitive ratio of any algorithm that solves the problem.

Hint: Theorems 2 and 3 of Lecture 3 notes.

2 Constraint Satisfaction Problems (15 points)

We term a search for a CSP solution *b-bounded* if for any level h in the search tree ($h < n$, the number of variables), we can push forward to find a value at level $h+1$ by considering possible values at level $h+1$ and at most reconsidering already chosen values at $b-1$ higher levels. Backtrack-free search is then the special case of 1-bounded search.

The width of a group of j consecutive vertices in an ordered constraint graph is defined as the number of vertices preceding the group in the ordering that are linked to any of the j vertices in the group. The j-width at a vertex v is the minimum, for $k = 1, \ldots, j$, of the width of the k consecutive vertices up to and including v. Thus the width at vertex v is the special case of 1-width. The j-width of an ordered constraint graph is the maximum j-width at all vertices in the ordering. The j-width of a constraint graph is the minimum j-width of all orderings of the constraint graph.

A CSP is (i,j)-consistent if, given values for i variables satisfying the constraints on those variables and given any other j (or fewer) variables, it will be possible to find values for those j variables such that the $i+j$ values taken together satisfy all constraints on the $i+j$ variables. Thus k-consistency is the special case of $(k-1,1)$-consistency. Strong (i,j)-consistency is (i',j)-consistency for all $i' \leq i$.

1. **(10 points)** Prove that, given a constraint graph for a CSP, there exists a vertical search order that guarantees j-bounded backtrack search if the problem is strongly (i,j)-consistent for i equal to the j-width of the graph.

   Hint: A generalization of the proof of Theorem 1 of Lecture 5 notes.

2. **(5 points)** Give a counterexample to the following claim. Let i, j, i', j' such that $i+j = i'+j'$; then a CSP is (i,j)-consistent if and only if it is (i',j')-consistent.

3 Social Choice (30 points)

Charles Dodgson, better known by his pen name Lewis Carroll, suggested the following voting rule in 1876. The *Dodgson score* of a given alternative $a \in A$, with respect to a given preference profile $\succ$, is the least number of exchanges between adjacent alternatives in $\succ$ needed to make a a Condorcet winner. For instance, let the set of voters be $N = \{1,2,3\}$, let the set of candidates be $A = \{a,b,c\}$, and let $\succ$ be given by:

1	2	3
a	b	a
b	a	c
c	c	b

In this example, the Dodgson score of a is 0 (a is a Condorcet winner), the score of b is 1, and the score of c is 3. The winner according to Dodgson's rule is the alternative with the *lowest* Dodgson score.

1. **(10 points)** It is known that manipulating Dodgson's rule is NP-hard. In lecture 7 we discussed an algorithm that efficiently decides the manipulation problem under any voting rule that satisfies certain conditions (Theorem 1 of Lecture 7 notes). Why can't we use this algorithm to efficiently manipulate Dodgson's rule? A proof is not needed, just an explanation.

2. **(20 points)** Consider the problem of computing the Dodgson score of a given alternative. The input is an alternative a and a preference profile $\succ$, and the output is the Dodgson score of a. Give an n-approximation algorithm for this problem, where n is the number of voters.

   Hint: Part 2 sort of contains a hint for part 1.

4 Cake Cutting (25 points)

Describe a discrete procedure for dividing a cake between 4 players that performs exactly three cuts and guarantees that each player believe they received at least 1/6 of the cake. Moving knives are not allowed and "marks" count as cuts.