CMPT 310
Artificial Intelligence Survey

Simon Fraser University
Summer 2011

Instructor: Oliver Schulte

Assignment 3: Chapters 6, 7, 8, 9. Logical Agents.

Total Marks: 100.
Due Date: ca. July 20, 10:20 am.

Instructions: Check the instructions in the syllabus. The university policy on academic dishonesty and plagiarism (cheating) will be taken very seriously in this course. *Everything submitted should be your own writing or coding.* You must not let other students copy your work. Discussion of the assignment is okay, for example to understand the concepts involved. If you work in a group, put down the name of all members of your group. On your assignment, put down your **name**, the number of the assignment and the number of the course. Spelling and grammar count.

Handing in the Assignment. Please post your assignment on our course management server https://courses.cs.sfu.ca/1114-cmpt-310-d100/ . Please include text answers and diagrams in a single pdf file. For code, XML files and such, provide a separate archive file. The whole package can be submitted as a .zip file. *The time when you upload your assignment is the official time stamp.*

We also need a printout. Please hand in the printout to the assignment box in CSIL (Computing Science Instructional Lab). *You need an access card for CSIL*. You should put the printout in the assignment box on the due date, but it doesn't have to be by 10:20 am.

Chapter 6, Constraint Satisfaction.

1. (35 + 20 bonus) Constraint satisfaction is closely related to logical inference. Let us explore this with the Wumpus world described in Section 7.2. At the same time, we'll start getting some idea about the problems of knowledge representation and knowledge engineering. Start by reading Section 7.2 about the Wumpus world. Let's simplify the Wumpus world as follows.

a. Let us focus only on the location of the Wumpus and the pit---ignore the gold, the walls, the arrow, moves by the agent etc.
b. Unlike the text we'll assume that the Wumpus cannot sit on a pit. So all location information is of the form "Wumpus is on square (x,y)" or "Pit is in square (x,y)" and these are mutually exclusive.
c. We assume that there are several Wumpuses, so any square can contain a Wumpus regardless of what the other squares contain.
d. We consider only the 9 squares in the 3x3 grid containing the initial square (1,1).
e. Since we are interested only in Wumpuses and the pit, the only relevant percepts in any given square are whether the agent feels a breeze or smells a stench. So in a square (x,y), the possibilities are (none, none), (stench, none), (none, breeze), (stench, breeze).

I. (11) Use the tool from AIspace http://www.aispace.org/constraint/index.shtml to draw a constraint graph with *binary constraints* that incorporates as many constraints as possible mentioned in the text. (Next to a Wumpus the agent smells a stench, next to a pit the agent feels a breeze, etc.). Note that there are no constraints between the content of a square (x,y) and the perception at the same square (x,y). You should be able to do this with 18 variables.
Note: It's often more convenient to edit the CSP text description rather than through the GUI. For this you can use the "View/Edit text description" button. Or you can open the .xml file in a text editor, which allows you to use the search and replace function. Even if you don't know XML, the CSP format is self-explanatory.

II. (2) How many variables would you need to extend this to the entire 4x4 grid?

III. (2) How many variables to extend it to a general *n*x*n* grid?

IV. (10) Verify your solution by using arc consistency to perform the following inferences. Show how the AIspace tool validates these inferences.
   a. If the agent's perception at square (1,1) is (none,none) (i.e., neither stench nor breeze), then there is neither a Wumpus nor a pit at square (1,2).
   b. If the agent feels no breeze at square (1,2), then there is no pit at square (1,3).

V. (10) As the text explains, the following inference is valid: Given that the agent perceives a breeze and no stench at square (2,1), and given that it perceives a stench and no breeze at square (1,2), it can infer that a Wumpus is in square (1,3). Show that although this inference is valid, it *cannot* be carried out in a graph with binary constraints between squares only. You can do that by producing a solution to your CSP graph that meets the following criteria.

a. The content of the square (1,1) is safe, neither a pit nor a Wumpus.
b. The perception at square (2,1) is (none,breeze).
c. The perception at square (1,2) is (stench,none).
d. There is no Wumpus in square (1,3).

You can find a solution meeting these criteria by using the tool (e.g. the Autosolve function), or by selecting values for variables yourself. Show the output of the tool that verifies your solution as consistent with all the constraints.

What to Submit for this question.

1. A screenshot of your constraint graph showing the nodes, domains, and constraints.
2. A listing of screenshots that show your various inferences. And CSP solutions.
3. Your .xml file with the CSP (loadable into the AIspace tool).

Bonus Question. (up to 20 points.) Show how you can extend the constraint graph specification to validate the conclusion that the Wumpus is in square (1,3). We will accept a precise outline of how you would go about this. An actual working implementation earns even more points. A few hints and suggestions.

I. Remember the option of editing the CSP in a text editor.

II. To cut down on the number of constraints you need, consider fixing the domains of the relevant variables, e.g. content of square(1,1) is “none”, perception at square (2,1) is (none,breeze), perception at square (1,2) is (stench,none). This is a bit of a hack because the CSP is then no longer a general model of the grid environment, but it could be easily done with a program for a general CSP.

Chapter 7. Propositional Logic and Resolution. (20 points total)

Resolution is a complete proof procedure for propositional logic, meaning that all logically valid conclusions ought to be provable using resolution. Let us check this claim in the Wumpus world. You can assume that we are in the 3x3 grid world. (In fact, the answers to the question are the same for the 3x3 and the 4x4 worlds.)

1. We use a variant of the notation given in Section 7.4.3. ("A simple knowledge base").
   P[x,y] is true if there is a pit in square [x,y].
   W[x,y] is true if there is a wumpus in [x,y].
   S[x,y] is true if the agent perceives a stench when in square [x,y].

2. Our knowledge base encodes the usual information about how the content of a square causes the perception at the square as discussed in the text. For instance, we have the biconditional

   S[1,1] iff (W[1,2] OR W[2,1])

   which is equivalent to the conjunction of the clauses

   not S[1,1] OR W[1,2] OR W[2,1]
   not W[1,2] OR S[1,1]
   not W[2,1] OR S[1,1]

I. (6) Which of these clauses are Horn clauses? How many Horn clauses do we have for each square [x,y]? You do not need to write out all the Horn clauses, just counting them "in your head" is sufficient.

II. (14) Suppose we add to our knowledge base that there is no pit at square [1,1] and there is no wumpus at square [1,1]. Draw a resolution tree that validates the inference "if the agent perceives a stench on square [1,2] and no stench at square [2,1], then there is a Wumpus at square [1,3]". That is, after adding the further premises

   {not P[1,1], not W[1,1], S[1,2], not S[2,1]}

   we can prove using resolution that

   W[1,3].

   A nice way to show your resolution proof is using a resolution graph as in Figure 7.13.

Chapter 8, 9. First-Order Logic. 35 points total + 10 points bonus.

I. (10 + 10 bonus) Represent the following sentences in first-order logic, using a consistent vocabulary (which you must define).

(a) (2) Every student who takes French passes it.

(b) (2) More than one student took French in Spring 2011.

(c) (2) Only one student took Greek in Spring 2011.

(d) (4) There is a male barber who shaves all men in town who do not shave themselves.

(e) Bonus question (10 points). There is a male barber who shaves all *and only* the men in town who do not shave themselves.

Is sentence (e) satisfiable?

II. (10 points) Go to AIspace and load the deduction tool from http://www.aispace.org/deduction/index.shtml. Load the sample knowledge base "Arithmetic Axiomatization".

1. (3) How can you ask the query "2+3=5?" in the vocabulary of this knowledge base?

2. (2) Use Depth-First Search (Deduction Algorithm) to prove this statement. Show the deduction tree with the order of goals proven.

3. (5) Use Breadth-First Search (Deduction Algorithm) to prove this statement. Show the deduction tree with the order of goals proven. Is there a difference with depth-first search? Briefly explain why there is/is not a difference.

III. (15 points) We now revisit the Wumpus world of the previous question in 1^{st}-order logic.

1. (10) Use the same vocabulary as above, but this time you may use 1^{st}-order variables. For instance, you can write W(X,Y) to express that there is a Wumpus in row X, column Y. Use the deduction tool from AIspace http://www.aispace.org/deduction/index.shtml to create a knowledge base of 1^{st}-order sentences that express the information about stenches and wumpuses in the Wumpus world. That is, replace the large set of biconditionals like

   S[1,1] iff (W[1,2] OR W[2,1])

   by *1st-order formulas*. Note: A small number of rules suffices.

   You may want to use the predicate **Adjacent** defined by the following facts: **Adjacent(1,2). Adjacent(2,1). Adjacent(2,3). Adjacent(3,2).**

   Create your knowledge base in the UBC tool. Verify the soundness of your knowledge base by validating the inference "if there is a Wumpus on square [1,3], then there is a stench on square [1,2]".

   What to Submit for this question. A loadable .pl file for your knowledge base in part III.1. A screenshot of your knowledge base and the proof trees for the valid inferences.