Institut for Matematik og Datalogi
Syddansk Universitet

December 4, 2009

Peter Schneider-Kamp

Introduction to Computer Science
E09 – Week 12

Announcements

The slides for the lectures on November 30 and December 2 are available from Blackboard. You will need these to solve the last assignment.

If you want to compare your results on the TSP assignment, you can find a list here: http://imada.sdu.dk/~marco/Teaching/Fall2009/DM526/results.txt

As mentioned in the last lecture, the code for the TSP must be updated in the file `Tour.java` as follows. Replace

```
public String toString() {
    String s = "";
    for (int i = 0; i < p.getSize(); i++) {
        s += (sequence[i]+1) + " " + p.getCoords(i)[0]+
                               " " + p.getCoords(i)[1]+"\n";
    }
    return s;
}
```

by

```
public String toString() {
    String s = "";
    for (int i = 0; i < p.getSize(); i++) {
        s += (sequence[i]+1) + " " + p.getCoords(sequence[i])[0]
                              +" " + p.getCoords(sequence[i])[1]+"\n";
    }
    return s;
}
```

Lecture: Monday, November 30

Marco Chiarandini started to lecture on artificial intelligence based on Chapter 11.3.

Lecture, Wednesday, December 2

Marco Chiarandini continued to lecture on artificial intelligence based on Chapters 11.1, 11.2, 11.6, and 11.7.

Lecture: Monday, December 7, 12-14 (U140)

Marco Chiarandini will start to lecture on artificial intelligence based on Chapters 11.4 and 11.5.

Discussion section: December 8, 10:15-12 (U9)

1. Exercise 3 on page 565 of text book.
2. Exercise 6 on page 566 of text book.
3. Exercise 47 on page 569 of text book.
4. Exercise 48 on page 569 of text book.
5. Which are the two properties of a heuristic to be useful in a search?
6. Indicate the size of the search space for at least one solution representation in the travelling salesman problem.
7. Fill in the blanks with input values that will cause the artificial neural network below to produce an output of 1.

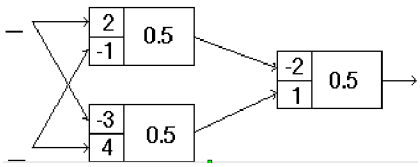

Figure 1: The neural network of Exercise 6.

8. The diagram in Figure 2 represents an associative memory as described in the text. If each unit has a threshold value of 0.5, what stable state will the system reach if it is initiated with the top and bottom units excited and the others inhibited? If the center unit has a threshold value of 2.5, under what condition will it become excited?

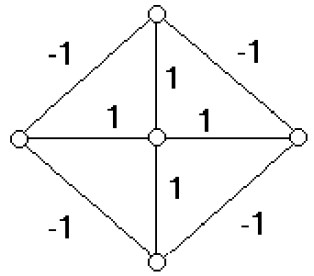

Figure 2: The associative memory diagram of Exercise 7.

9. Which of the conlcusions below would require the closed-world assumption in the context of a database that contained a list of subscribers to the New York Times.

   (a) John Doe subscribes to the New York Times.

   (b) John Doe does not subscribe to the New York Times.

   (c) Either Mary Doe or John Doe does not subscribe to the New York Times.

   (d) Either Mary Doe or John Doe subscribes to the New York Times.

10. In the following lists, match the terms indicated by a the lower case letter with their corresponding phrase indicated by an upper case letter.

(a) real-world knowledge	(g) Turing test
(b) heuristic	(h) associative memory
(c) best first	(i) expert system
(d) genetic algorithms	(j) semantic net
(e) inference rule	(k) artificial neural network
(f) uniform-cost search	(l) state graph

(m) image analysis

(n) contextual analysis

(o) linguistics

(p) perceptron

(q) greedy heuristic

(r) syntax analysis

(s) semantic analysis

(t) agent

(u) breadth-first search

A A responsive entity

B A means of measuring a machines ability to perform like a human

C The task of understanding an image

D A rule of thumb that leads to acceptable solutions but for which there is no formal proof of its correctness.

E The result of considering all options equally important

F A breadth first search in which nodes with best evaluation are expanded first

G A single descent in a search tree dictated by heuristic decisions.

H A breadth first search in which nodes with lowest path cost are expanded first

I A picture of all states and productions

J A means of obtaining a statement that is a logical consequence of other statements

K The database used by an intelligent system to support its reasoning

L A multiprocessor computer consisting of many simple processors

M A field of artificial intelligence that applies evolutionary theories to the software development process

N The ability to recall related information

O A software package for solving problems within a particular field

P A means of representing knowledge

Q To relate a sentence to its environment

R The study of human communication

S A single layer neural network for learning linear separable functions.

T Parsing a sentence identifying the grammatical role of its parts

U Identifying the meaning of words in a phrase.

11. If you were the interviewer in a Turing Test how would you plan your conversation with the human and the machine in order to distinguish who is who? Try your conversation with ALICE (link) and report back the discussion. Comment the answers emphasizing the elements that let you understand you are talking with a machine and not an human being. Make up possible explanations of the functioning of ALICE.

Prepare the discussion on the following issues. In the exercise session you will be covering a role and you will have to argue from the perspective of that role. Hence it is important that in the preparation for each issue you think about arguments as well as counter arguments.

1. Social Issue. N. 2 and 3.
2. Social issue N. 7.
3. Social issue N. 10.
4. Imagine that in two years from now you sit in the student council of the faculty. The subject under discussion at the next meeting is the introduction of a new regulation at SDU that prohibit the use of the computer during lectures unless expressedly requested by the lecturer. Prepare arguments both in favor and against this decision. At the exercise session you will discuss this with a colleague and you will have to cover one of the two roles. Hence, it is important that during the preparation you try to foresee possible criticism to your arguments and prepare justifications.
5. In the lecture on AI you have been presented with *functionalism*, a theory of mind. Accordingly, mental states are identified by a functional role and are able to be manifested in various systems, even perhaps computers, so long as the system performs the appropriate functions. While computers are physical devices with electronic substrate that perform computations on inputs to give outputs, so brains are physical devices with neural substrate that perform computations on inputs which produce behaviors.

   Important propositions in the philosophy of AI sustaining this position include:

- Turing's "polite convention": If a machine acts as intelligently as a human being, then it is as intelligent as a human being.
- The Dartmouth proposal: "Every aspect of learning or any other feature of intelligence can be so precisely described that a machine can be made to simulate it."
- Newell and Simon's physical symbol system hypothesis: "A physical symbol system has the necessary and sufficient means of general intelligent action."
- Searle's strong AI hypothesis: "The appropriately programmed computer with the right inputs and outputs would thereby have a mind in exactly the same sense human beings have minds."
- Hobbes' mechanism: "Reason is nothing but reckoning."

Make a list of advantages of functionalism and a list of arguments against it. Be prepared to cover both roles of proponent and opponent of this account of mind in the exercise session and to justify your corresponding arguments.

6. The following is the description given by Wikipedia of Searle's *Chinese room* thought experiment[1]

> Searle's thought experiment begins with this hypothetical premise: suppose that artificial intelligence research has succeeded in constructing a computer that behaves as if it understands Chinese. It takes Chinese characters as input and, by following the instructions of a computer program, produces other Chinese characters, which it presents as output. Suppose, says Searle, that this computer performs its task so convincingly that it comfortably passes the Turing test: it convinces a human Chinese speaker that the program is itself a human Chinese speaker. To all of the questions that the human asks, it makes appropriate responses, such that any Chinese speaker would be convinced that he or she is talking to another Chinese-speaking human being.

[1] A thought experiment is used to test or illuminate a hypothesis or theory even though it may not be possible to actually perform it.

> Some proponents of artificial intelligence would conclude that the computer "understands" Chinese. [...]
>
> Searle then asks the reader to suppose that he is in a closed room and that he has a book with an English version of the aforementioned computer program, along with sufficient paper, pencils, erasers and filing cabinets. He can receive Chinese characters (perhaps through a slot in the door), process them according to the program's instructions, and produce Chinese characters as output. As the computer had passed the Turing test this way, it is fair, says Searle, to deduce that the human operator will be able to do so as well, simply by running the program manually.

What is the target Searle wishes to achieve with this thought experiment? Is it supporting or rejecting the strong AI position?

7. A slightly different, although related account of mind is *connectionism*, Accordingly, mental or behavioral phenomena are modelled as emergent processes of interconnected networks of simple units. More broadly emergent behavior of fundamental particles is a contemporaneous theory to describe complex systems like economics, fluid dynamics and swarms of insects. Scientific holism holds that the behavior of a system cannot be perfectly predicted by its elements alone, no matter how much data is available. Natural systems can produce surprisingly unexpected behavior, and it is suspected that behavior of such systems might be computationally irreducible, which means it would not be possible to even approximate the system state without a full simulation of all the events occurring in the system. Stephen Wolfram in his controversial "A New Kind of Science" argues that, rather than traditional mathematics, computational systems and fundamental laws, which can be described as simple programs, are needed to model and understand complexity in nature.

   Try to formulate and justify criticism against these positions. If you embrace them, try imagine the criticism of opponents and your possible answers.

Lab: December 11, 10:15-12 (terminal room above U49)

In this session you will be asked to implement a local search algorithm for the TSP. The session will be carried out in the form of a competition. You will work in pairs and try to achieve the best results on the instances provided. In preparation, make sure that you have all the code delivered with the TSP Assignment installed on the machine that you will use together with the instance data. Bring also your implementation of the construction heuristic that was asked for the Assignment. The instructor will guide you throughout the details of the session.

Assignment due 14:15, December 17

Late assignments will not be accepted. Working together is not allowed. You may write this either in English or Danish. Write clearly if you do it by hand. Even better, use LaTeX.

Heuristic Search and Neural Networks

1. Given the hypothetical search tree of Figure 3, List the order in which the nodes will be visited in a (i) breadth-first search, (ii) depth-first search, (iii) uniform cost search, (iv) best-first search, (v) greedy construction. The values on the edges of the tree represent the step cost while the h values on the nodes next to the letter that represent the node are the heuristic values of completing a solution from that node. [Hand in electronically in a text file with name ex1.txt and write your answers to the 5 points in 5 separated lines. Each line is a string of uppercase letters without space representing the visit order.]

2. Consider the application of a perceptron network in classification. For example, consider the application in medical diagnosis of diabetes. The features in the pattern space may be blood sugar measurement and amount of fluid intake per day of different patients. These features are used to classify the symptoms sets into two different diagnosis categories, that is, presence or not of diabetes.

   Table 1 reproduces some past measurements with the relative category. Here +1 means that the person had diabetes, while -1 that he/she did not. (The data are available also online: link)

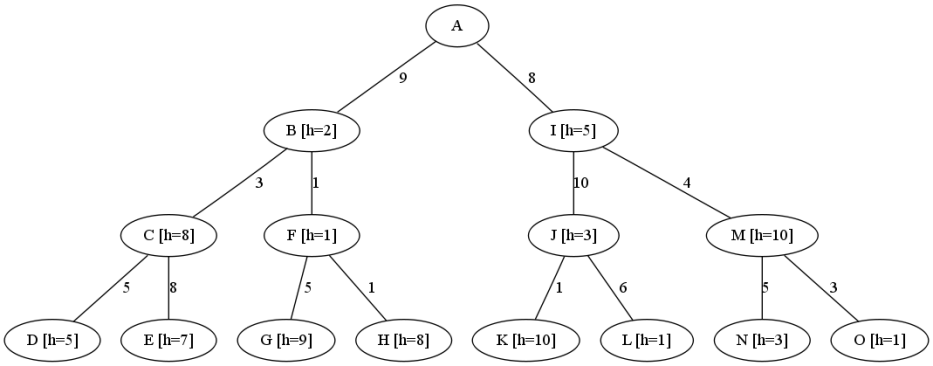

Figure 3: The search tree for exercise 1. Values on the edges represent step costs and values on the nodes the heuristically computed value of completing the solution from that node.

Apply the Perceptron Training Algorithm to learn the weights for a perceptron with x_1,x_2 and *bias* as inputs and $-1, +1$ as output.

You can carry out the computations using `maple` or manually (in this latter case 5 iterations are sufficient). Report the final weights and a plot of the points together with the linear separator provided by the perceptron.

x_1	x_2	Diabetes
1.0	1.0	1
9.4	6.4	-1
2.5	2.1	1
8.0	7.7	-1
0.5	2.2	1
7.9	8.4	-1
7.0	7.0	-1
2.8	0.8	1
1.2	3.0	1
7.8	6.1	-1

Table 1: The data for Exercise 2.

3. Consider Neural Networks with inputs in the range $[0, 1]$ and with a step function g. A network is defined by the weights on the links and a threshold value of g at each node.

(a) In Boolean logic, the majority function is a function from n inputs to one output. The value of the operation is false when $n/2$ or more arguments are false, and true otherwise. Draw a network that represent the majority function for 4 input nodes.

(b) Draw a network that represent the "exactly two out of three" function for three inputs.

(c) Draw a network to simulate the XOR operator in Boolean logic. XOR (exclusive-or) is a logical operator that results in the output being true if one of the inputs, but not both, is true. If both inputs are true the output is false.