

You are allowed to have a two-sided A4-sized hand-written “cheat sheet” with your own notes written on it. You are also allowed to have a calculator (but it is not required). Note that the cheat sheet must be returned with you answers. (You can get it back later.)

Write on each answer sheet: the name of the course, the date, your student number, name, and signature. Please use **clear handwriting** in your answers.

You can answer in Finnish, Swedish, or English.

NB: Please check that you received both problem sheets. Answer four (4) problems out of five (5). If you answer all problems, only the first four will be taken into account. Each problem is worth max 10 points.

1. Philosophy and History of AI

- a. (6 p) Essay. A good length is about 2/3 of a page.

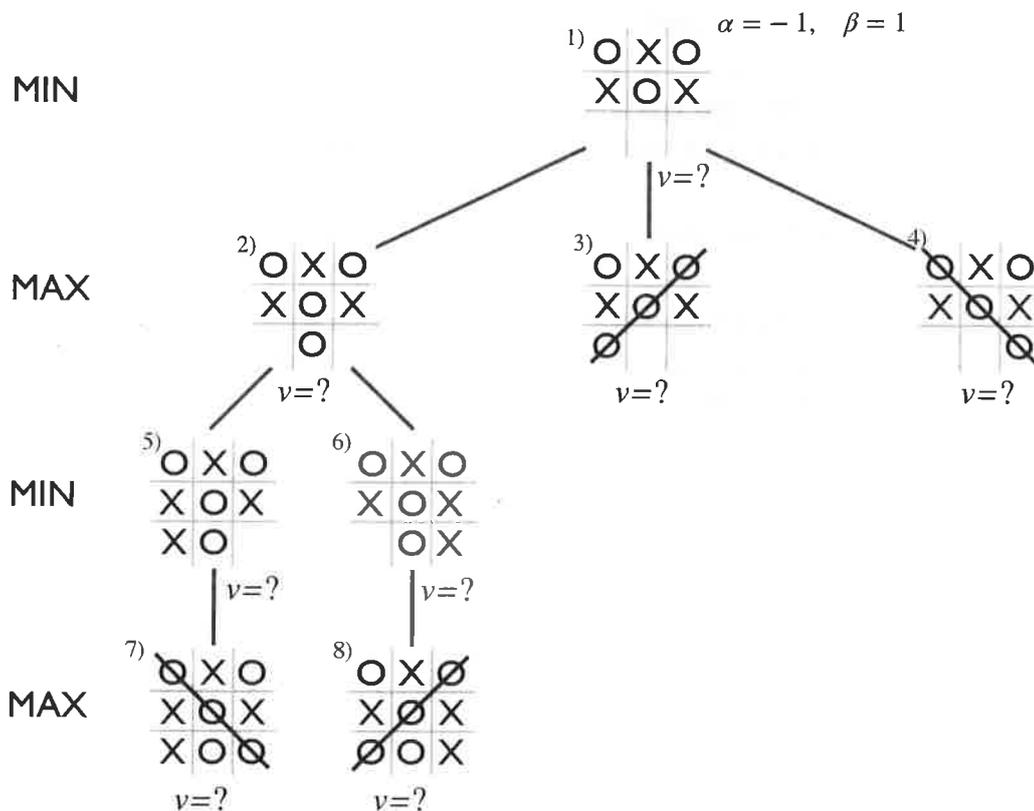
Topic: *The first thing that comes to mind when we say "AI"? Today and in the past*

Imagine you'd time travel back in time to another decade. What would the AI experts of that time think when you say "AI"? Time travel to a couple different decades and answer the above question.

- b. (2 p) Is the claim “Deep learning solves everything” true or not? Give a brief justification to your answer.
- c. (2 p) Explain the Turing test. What does it measure? What does it *not* measure?

Search and Games

Consider the tic-tac-toe position shown at the root node (1) of the below game tree. We've drawn out the game tree for you. (Isn't that helpful!)



- (2 p) Fill in the values of the nodes. The value of any state where the Min player who plays with the noughts 'O' wins is -1 (so for example, the value at node (4) is -1), and the value of any state where the Max player (the crosses, 'X') wins is $+1$. If the game ends in a tie, the value is 0 .
- (1 p) What is the value of (the position at) the root node? What does that mean?
- (5 p) If we alpha-beta pruning at the root node with initial values $\alpha = -1$, and $\beta = 1$, and the children of each node are processed from left to right, what pruning takes place, if any? Explain your answer using the α and β values and their updates.
- (2 p) Is alpha-beta pruning useful in games like chess or Go? Is it enough?

3. Natural Language Processing

Consider a context-free grammar with the following rules:

$$\begin{array}{lll}
 S \rightarrow S S & T \rightarrow S R & R \rightarrow \text{right} \\
 S \rightarrow L T & L \rightarrow \text{left} & S \rightarrow \text{dot}
 \end{array}$$

S is the starting symbol. From S, we can generate, for example, $S \rightarrow SS \rightarrow LTS \rightarrow LTLT \rightarrow LSRLT \rightarrow LSRLSR \rightarrow \dots \rightarrow \text{left dot right left dot right}$.

- (4 p) Does the language contain the following three sentences (or strings): "" (the empty string), "dot dot dot", "left dot right left dot right"? Justify your answer.
- (4 p) Apply the CYK algorithm to fill in the following table in order to determine whether the sentence "left left dot right right" belongs to the language. *Hint*: Cells (4,4), (5,5), and (4,5) have already been filled in as much as they ever will. Recall that for example, the cell (4,5) covers the words 4–5: "right right".

(1,5)					
(1,4)	(2,5)				
(1,3)	(2,4)	(3,5)			
(1,2)	(2,3)	(3,4)	(4,5)		
(1,1)	(2,2)	(3,3)	R (4,4)	R (5,5)	
<i>left</i>	<i>left</i>	<i>dot</i>	<i>right</i>	<i>right</i>	

- (2 p) What is the corresponding parse tree (or parse trees), if any? Does the sentence belong to the language or not?

Reasoning under Uncertainty

- a. (3 p) What is the Bayes rule? Give an example of a concrete application (with real-world events, not just A, B, \dots). Use meaningful numerical values, none of which is 0.0 or 1.0, and apply the Bayes rule to them.
- b. (2 p) Explain the importance of the Bayes rule in AI. *Hint*: The question isn't asking what makes probabilistic inference so important.
- c. (3 p) How do you generate data from a given Bayesian network? Explain the procedure step-by-step.
- d. (2 p) How can you estimate the probability $P(\text{"It will snow next weekend"} \mid \text{"The temperature on Tuesday was } 7^\circ \text{ Celsius"}, \text{"It is October"})$, when you have generated 100 000 tuples (instances) of data from a Bayesian network including the above events?

Neural Networks

- a. (4 p) Be a professor! Come up with neural network related **questions** when the correct answers are below. The answer should be easy to guess given the question (by someone who has taken the Intro to AI course).
 - i. a neural network model where one neuron wins when an input is presented
 - ii. linearly separable data
 - iii. deep learning
 - iv. activation function
- b. (1 p) Explain the basic idea of learning neural networks. Most importantly, what changes in the process of learning?
- c. (4 p) Describe the operation of a kind of neural network of your choice (for example, a multilayer perceptron): what is the structure of the input data, and what happens when an input is presented? NOTE: The network should have more than one neuron.
- d. (1 p) What would be a typical application of this kind of a neural network?