You can have with you in the exam a two-sided A4-sized "cheat sheet" that has to be turned in together with your answers. Note that every student must return at least an empty answer sheet with their student number and name.

Please write the name of the course, the data, your student number, name, and signature in each answer sheet.

NB: Check that you have all the exam sheets (6 questions). Answer five (5) questions out of six (6). In case you answer all six, only the first five will be taken into account in grading. The maximum points for each question is eight.

1. **Philosophy of AI, etc.**

a. (2 p) Explain Turing test. (NB: not Turing machine.) What kind of artificial intelligence (AI) has been reached when the Turing test can be passed.

b. (2 p) In movies, AI often gains consciousness and attacks humans. Mr Dionarap is worried about this risk and campaigns for a law against all AI algorithms. What practical consequences would such a law have? Give a couple of examples. What problems do you see in enforcing the law?

c. (2 p) Ms Noruen claims that AI methods cannot be implemented on conventional computing hardware. She thinks that only special computational models based on neural networks, where computation is parallel and asynchronous, are needed for the job. Do you think that Ms Noruen is right? Justify your answer.

d. (2 p) What kind of problems were faced in the 80s when AI problems were being tried to solve using logic? What was the reaction on these problems?

2. **Search as problem solving**

The task is to solve a puzzle where the starting state A is on the left, and the goal state M is on the right. Allowed moves between states are such that a number that is next to (on the left, right, above or below) an empty slot moves to the empty slot. For instance, both states 1 and 2 below can be reached directly from the starting state A:

```
   +-+-+      +-+-+       +-+-+     +-+-+
   |2|3|      |1|2|       |2|3|     | |3|
A: +-+-+   M: +-+-+    1: +-+-+  2: +-+-+
   | |1|      |3| |       |1| |     |2|1|
   +-+-+      +-+-+       +-+-+     +-+-+
```

a. (2 p) Give all the states that can be reached from the starting state as a state diagram. Also indicate the allowed moves between the states.

b. (3 p) Simulate both depth-first and breadth-first search in the state diagram of item a. (It is sufficient to list the states in the order in which they are visited.) How many states are visited before each search method finds the goal state? Which one finds a better solution (less moves)? Is the result always the same no matter in which order the states are expanded?

c. (3 p) Let's assume that you are going to implement a TravelPlanner (Reittiopas). Given any two public transportation stops A and B, the day of the week and time of day, the TravelPlanner finds the fastest route from A to B using information about public transportation routes and timetables.

   How would you set off solving the task? How would you represent the "state" of the search, and what would be the possible transitions between the states? What kind of a search algorithm would be suitable? Would you use a heuristic, and if yes, which one?

3. **Image Processsing**

a. (3 p) Why do digital signals such as image and sound cause problems for logic-based and other GOFAI ("Good Old Fashioned AI") methods? What kind of methods are more suitable for handling digital signals and why?

b. (3 p) If the task were to implement a face recognition method that can identify a person using a camera, what kind of method would you choose? Briefly explain how it works. Would it always work or would it have some particular restrictions?

c. (2 p) Below is a pseudocode of a denoising algorithm. Briefly explain how it works and what is the idea underlying it.

```
denoising(x,t):
    n ← length(x)
    c ← dwt(x)          // discrete wavelet transform
    for i = 1,...,n:
        if |c[i]| < t: c[i] = 0
    end-for
    x ← idwt(c)         // inverse wavelet transform
    return x
```

4. **Neural Networks**

a. (2 p) True of false?
   i. The neurons in a multi-layer perceptron compete about which one of them gets to activates.
   ii. A recurrent ("takaisinkytkeytyvä") neural network can be used to implement a fault-tolerant memory.
   iii. A Bayesian network is an example of a stochastic (i.e., one that is based on randomness) neural network.
   iv. For each neuron in a neural network, there is an associated set of weights that together with the inputs of the neuron, determine the activation.

b. (2 p) Describe the basic principle of a neural network of your choice (*). What are the inputs of the network and what happens when an input is given to the network?

c. (2 p) Briefly explain the learning method of the network you described in item b. The basic idea is enough, the details are not necessary. (Do not give pseudocode, but instead, explain in words.) What kind of training data is used?

d. (2 p) What would be a typical application of the network?

*) NB: The network has to have more than one neuron.

5. **Probabilistic modeling**

Here are some probability formulas to refresh your memory:

(i) $P(\neg A) = 1 - P(A)$ "negation"

(ii) $P(A \vee B) = P(A) + P(B) - P(A, B)$ "disjunction"

(iii) $P(A) = P(A, B) + P(A, \neg B)$ "marginalization"

(iv) $P(A_1,\ldots,A_k) = P(A_1)\, P(A_2 \mid A_1)\, P(A_3 \mid A_1,A_2) \ldots P(A_k \mid A_1,\ldots,A_{k-1})$ "chain rule"

(v) $P(A \mid B) = P(A)\, P(B \mid A) / P(B)$ "Bayes rule"

(vi) $A \perp B \mid C \Rightarrow P(A,B \mid C) = P(A \mid C)\, P(B \mid C)$ "conditional independence"

a. (2 p) Let $P(V \mid S) = 0.25$ ja $P(V \mid \neg S) = 0.0005$, where S means that a message is junk mail (spam) and V means that the word "Viagra" appears in the message. Evaluate the probability $P(S \mid V)$, when $P(S) = 0.2$. You can give the answer as a rational number, i.e, in the form r/s, where both r and s are integer (kokonaislukuja).

b. (2 p) The following table contains the number of occurrences of certain words in both spam and ham (normal mail).

word	spam	ham
money	24	6
lottery	16	4
Bayes	1	10
total count	2000	10000

Estimate probabilities $P(\text{Word}_i = \textit{money} \mid \text{spam})$, $P(\text{Word}_i = \textit{lottery} \mid \text{spam})$, and $P(\text{Word}_i = \textit{Bayes} \mid \text{spam})$ as well as the corresponding probabilities for ham.

c. (4 p) Compute the probability that a message is spam when it says "*Bayes lottery money*". Use a naive Bayes model where the words are independent of each other given the type of the message (spam/ham). Use the prior $P(\text{spam}) = 0.2$. You can again give the answer as a rational number.

6. **Robotics**

Because building Legos is fun, in this question you get to design your own Lego robot and its program!

Below are pictures of available parts for the robot:

- a central unit with sockets for two motors (M1,M2) and three sensors (S1,S2,S3),
- three wheels,
- two motors (they can be controlled using commands *forward(), backward(), stop()*),
- an ultrasonic sensor (you can get the distance to the nearest obstacle in cm by calling method *getDistance()*),
- a light sensor (you can get a value describing the "lightness" of the floor under the sensor by calling method *getLightValue()*)

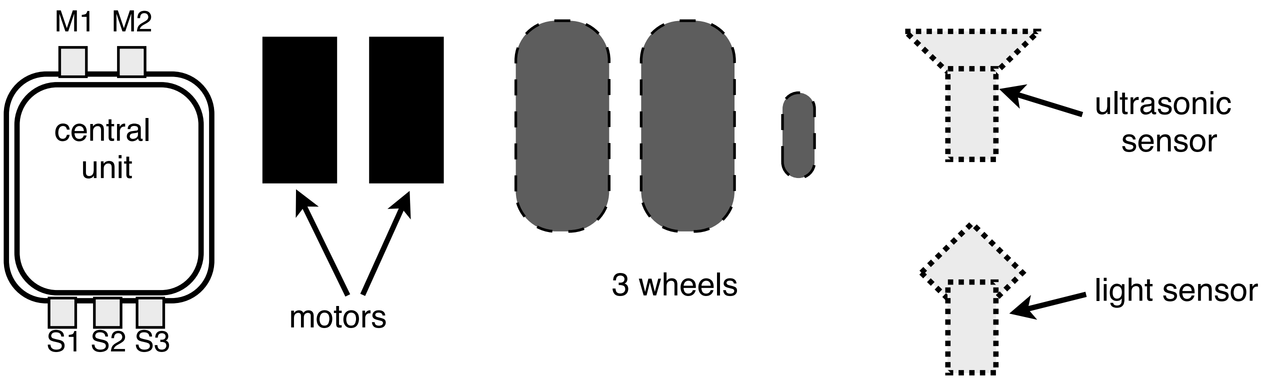

In addition, you can of course use any regular Lego blocks, axes, cables, and other parts you may need.

a. (2 p) Draw a simple robot using which you can solve the task in the next item. Include the cables connecting the right components to the central unit in the drawing.

b. (4 p) Give pseudocode of a program that follows a black line drawn on the floor until it arrives at an obstacle that it can observe using the ultrasonic sensor. Once the robot observes the obstacle, it should stop. You can assume that the robot is initially placed on top of the black line.

   Hints: The command *M1.forward()* makes the motor connected to socket M1 start moving (forward). You can get inputs from the sensors by calling *S1.getDistance()* etc. You don't need to give the constructors or other initializations of the motors and the sensors.

c. (2 p) What does the program planned on a piece of paper (such as now) usually fail on the first try? Mention some features of the physical environment that cause additional problems compared to "normal" (non-robotic) programming and describe how they could be solved.