Artificial Intelligence: Exercises for Tutorial 1 on Propositional Logic

October 30, 2021

Introduction

The following multiple choice questions are examples of typical questions one can expect on the AI exam. The questions on the AI exam are also multiple choice, but for this tutorial one has to explain the answers given. Moreover at the end one can find some open questions. After the tutorial the answers to the MC will be available on Canvas.

Exercises

1. Consider the proposition:

   $$R \Rightarrow (\neg R \Rightarrow W)$$

   How many models are there for this proposition? (Do model checking, make a truth table.)

   (a) 2

   (b) 4

   (c) 6

   (d) 8

2. Consider the proposition:

   $$R \Rightarrow (\neg S \Rightarrow W)$$

   How many models are there for this proposition? (Do model checking, make a truth table.)

   (a) 1

(b) 3

(c) 5

(d) 7

3. We are given the following premises:

   1. $bread \vee earlyMeeting$
   2. $(tea \vee \text{coffee}) \wedge juice$
   3. $earlyMeeting \Rightarrow yoghurt$
   4. $yoghurt \Rightarrow \neg coffee$
   5. $\neg yoghurt$

   The question is whether we can prove *bread* from these premises using resolution. Which of the following answers is correct?

   (a) Yes, the conclusion follows

   (b) No, the conclusion does not follow, but if you add the premise *juice* the conclusion can be derived

   (c) No, the conclusion does not follow, but if you add the premise *tea* the conclusion can be derived

4. We are given the following premises:

   - $(P \vee Q) \wedge (P \vee T)$
   - $(Q \wedge T) \Rightarrow (V \Rightarrow W)$
   - $\neg[(T \Rightarrow S) \Rightarrow \neg(S \Rightarrow W)]$

   The question is whether we can prove $V \Rightarrow S$ from these premises. Which of the following answers is correct?

   (a) Yes, the conclusion follows.

   (b) No, the conclusion does not follow, but if you add the premise T the conclusion can be derived.

   (c) No, the conclusion does not follow, but if you add the premise $\neg S$ the conclusion can be derived.

   (d) No, the conclusion does not follow, but if you add the premise V the conclusion can be derived.

   The way to solve this type of problems is as follows. Apply resolution either until you have found a contradiction (and then alternative (a) is apparently correct) or until there is no valid application of the resolution rule that produces a new sentence. In the latter case, comparing the remaining three alternatives with the list of sentences you have constructed by means of resolution will quickly tell you the correct alternative.

5. (For discussion. There is not one correct answer, I guess.)
   In order to let a machine reason logically it is required to use a language that has a clear semantics, i.e. names denote one single object and predicates and function symbols have unambiguous meanings. Do you think this makes a machine a better reasoner than people? What about vague predicates like "young", "large", "friendly", should we try to avoid them to become more "logical" and "rational" or is it possible to make the machine also be competent to use these fuzzy terms in a "logical" way?