

# UG Exam Performance Feedback

## Second Year

### 2016/2017 Semester 2

---

COMP24412 Symbolic AI

Ian Pratt-Hartmann

Comments Notes on Comp24412 Exam (May, 2017)

General remarks: Performance was well below the standard of previous years. Many Candidates had quite evidently not read the course textbook, not thought about what they had done in the laboratory exercises, and not really mastered the relevant theory, often to the extent that they could not formalize simple English sentences in first-order logic. In addition, far too many Candidates were hampered by appalling examination technique, particularly in respect of the ability to give clear definitions in grammatical---or even just comprehensible--English.

Q1. Performance on this straightforward question was particularly poor. Many Candidates could not articulate the answers to 1a and 1d properly. Answers to 1c and 1e were frequently incorrect. Worst of all, 1g was answered correctly by only a select few: the question clearly states that non-ground terms should not be unified. Many Candidates demonstrated in this part a complete inability to write Prolog programs. Parts 1f and (to my surprise) 1h were generally correctly answered.

Q2. Few Candidates attempted this question---fortunately. Part a) requires some simple list processing to be performed in Prolog: almost no Candidates realized that exceptions need to be made to take account of the fact that the table never gets full when a block is moved to it; in fact, almost no one got close. Parts c) and d) were also poorly attempted. A small number of Candidates, who had learned the material in the lecture slides, got very high marks, showing that there was nothing intrinsically wrong with the question. The last part required Candidates to realize that theorem provers which prove existentially quantified formulas (in this case, a formula asserting the existence of a plan) can be made to return witnessing objects (in this case, a plan). That defeated everyone, I think.

Q3. This was slightly better done. At least a number of Candidates knew the answer to 3a could perform the routine lambda-calculus calculations required in 3b. Many Candidates did not read question 3b properly; this clearly states that they were to derive the meaning of the sentence, not its phrase-structure alone. Many Candidates made mistakes in 3c (so simple). Answers to 3d were so-so.

Q4. This was probably the best done question (or perhaps the most leniently marked). Even so, when I set it, I had no idea that such a large proportion of the class would struggle with formalizing the simple sentences in 4a. Those who get this wrong presented the Marker with a quandary. If subsequent calculations are done correctly (and they often were), Candidates should not lose marks again (and indeed they did not). However, in some cases, errors made in 4a made the subsequent theorem proving task in 4d easier than it would have been. A small number of marks were taken off in this case only if errors in 1a so trivialized the proof that it was obviously wrong. Resolution proofs were in many cases very poorly set out, so that it was impossible to see what was supposed to be derived from what.

---