Answers

The University of Nottingham

SCHOOL OF COMPUTER SCIENCE

A LEVEL 3 MODULE, AUTUMN SEMESTER 2017-2018

KNOWLEDGE REPRESENTATION AND REASONING

Time allowed TWO hours

Candidates may complete the front cover of their answer book and sign their desk card but must NOT write anything else until the start of the examination period is announced

Answer all FOUR questions

Only silent, self contained calculators with a Single-Line Display are permitted in this examination.

Dictionaries are not allowed with one exception. Those whose first language is not English may use a standard translation dictionary to translate between that language and English provided that neither language is the subject of this examination. Subject specific translation dictionaries are not permitted.

No electronic devices capable of storing and retrieving text, including electronic dictionaries, may be used.

DO NOT turn your examination paper over until instructed to do so

1. This question is on expressing knowledge in first-order logic and logical entailment.

   (a) Translate the sentences below from English into first-order logic. (10 marks)

   Use constants m for Colonel Mustard, p for Mrs Peacock, l for library, and t for 12 p.m. Please use predicates $Murderer^4$, In^3, O^1, M^1, where $Murderer(x, y, z, u)$ stands for x murdered y at place z at time u, $In(x, y, z)$ stands for x was at place y at time z, $O(x)$ means x had an opportunity and $M(x)$ means that x had a motive.

   S1 Somebody murdered Colonel Mustard in the library at 12 p.m., and there is exactly one person who is the murderer of Colonel Mustard.

   S2 Anybody who was in the library at 12 p.m. had an opportunity to murder Colonel Mustard.

   S3 Mrs Peacock was in the library at 12 p.m.

   S4 More than one person had a motive to murder Colonel Mustard.

   S5 A person is a murderer of Colonel Mustard if only if they had a motive to murder Colonel Mustard and an opportunity to do so.

   Answer. **A** (where **K** is Knowledge, **C** Comprehension, **A** Application)

   S1 $\exists x(Murderer(x, m, l, t) \wedge \forall y(Murderer(y, m, l, t) \supset y = x))$

   S2 $\forall x(In(x, l, t) \supset O(x))$

   S3 $In(p, l, t)$

   S4 $\exists x \exists y(M(x) \wedge M(y) \wedge \neg(x = y))$

   S5 $\forall x(Murderer(x, m, l, t) \equiv M(x) \wedge O(x))$

   (b) Give a definition of logical entailment. (5 marks)

   Answer. **K**

   A set of sentences S logically entails a sentence α iff in every interpretation where S are true, α is also true.

   (c) Do the sentences from part (a) logically entail that Mrs Peacock murdered Colonel Mustard? Justify your answer referring to the definition of logical entailment. (10 marks)

   Answer. **A**

   It is possible to construct an interpretation where S1-S5 are true and $Murderer(p, m, l, t)$ is false.

   In the interpretation, Mrs Peacock does not have a motive, and another person was in the library at 12 p.m. (another person has an opportunity). Formally, $J = (I, D)$, the domain D contains objects $d_1, \ldots, d_4$ denoted by m, p, l, t, and in addition d_5 and d_6 such that $\langle d_5, d_1, d_3, d_4 \rangle \in I(Murderer)$ (d_5 is the murderer). $I(Murderer)$ does not contain any other tuples, so S1 is true.

   $I(In) = \{\langle d_5, d_3, d_4 \rangle, \langle d_2, d_3, d_4 \rangle, \langle d_1, d_3, d_4 \rangle\}$. (We don't have to add that Colonel Mustard was also in the library at 12 p.m., but I added it for extra realism.) In particular, Mrs Peacock is in the library at 12 p.m. so S3 is true.

   $I(O) = \{d_1, d_2, d_5\}$ (so S2 is true).

   $I(M) = \{d_6, d_5\}$. So S4 is true.

Now the only person who is in $I(O)$ and $I(M)$ is d_5, who is also the murderer. S5 is true. All of S1-S5 are true and $Murderer(p, m, l, t)$ is false.

(4 marks for proposing to find an interpretation where S1-S5 are true and $Murderer(p, m, l, t)$ is false. Further 6 marks allocated pro-rata depending on whether the interpretation indeed does that, and the answer justifies that it does.)

2. This question is on clausal form, resolution and unification.

(a) Reduce the following sentences to clausal form: (10 marks)

S1 $\forall x \forall y \forall z (R(x,y) \land R(x,z) \supset \exists u (R(y,u) \land R(z,u)))$

S2 $\forall x \forall y \forall z (R(x,y) \land R(y,z) \supset R(x,z))$

S3 $\forall x R(x,x)$

S4 $\forall x \forall y \forall z (R(x,y) \land R(x,z) \supset R(y,z))$

S5 $\forall x \forall y (R(x,y) \supset R(y,x))$

Answer. **A** (2 marks per sentence)

S1 $[\neg R(x,y), \neg R(x,z), R(y, f(x,y,z))], [\neg R(x,y), \neg R(x,z), R(z, f(x,y,z))]$

S2 $[\neg R(x_2,y_2), \neg R(y_2,z_2), R(x_2,z_2)]$

S3 $[R(x_3,x_3)]$

S4 $[\neg R(x_4,y_4), \neg R(x_4,z_4), R(y_4,z_4)]$

S5 $[\neg R(x_5,y_5), R(y_5,x_5)]$

(b) Show by resolution that S1-S4 above logically entail S5. (10 marks)

Answer. **A**

The students first need to produce clausal form of negation of S5:
$\neg \forall x \forall y (R(x,y) \supset R(y,x)) = \exists x \exists y (R(x,y) \land R(y,x))$
which become

S5a $[R(a,b)]$

S5b $[\neg R(b,a)]$

(5 marks for doing this correctly, 3 marks for thinking about negation of S5 even if they don't get it right).

The simplest derivation (obviously students may come up with something more complicated but still correct; 5 marks for a correct derivation)

i. $[\neg R(a,z_4), R(b,z_4)]$ from S5a, S4, x_4/a, y_4/b
ii. $[R(b,a)]$ from i, S3, x_3/a, z_4/a
iii. [ ] from ii, S5b

(c) For the pairs of literals below, state whether they unify, and if yes give a most general unifier for them. Note that x, y, z, u, w are variables and a, b constants.

i. $P(f(x), g(a), h(x,y))$ and $P(f(g(b)), g(z), u)$ (2 marks)
ii. $P(f(x), x, h(x,y))$ and $P(z, g(z), h(u,w))$ (2 marks)
iii. $P(f(a), g(a), h(x,y))$ and $P(f(g(b)), g(z), u)$ (1 marks)

Answer. **A**

i. $P(f(x), g(a), h(x,y))$ and $P(f(g(b)), g(z), u)$ unify with mgu $\{x/g(b), z/a, u/h(g(b),y)\}$
ii. no; impossible to make $f(x), x$ and $z, g(z)$ the same by any substitution.
iii. no: there is no free variable in the first two arguments, $f(a)$ and $f(g(b))$

3. This question is on Horn clauses, rule-based systems and propositional resolution.

(a) Define Horn clauses and give examples of positive Horn clauses and negative Horn clauses. (5 marks)

Answer. **K,C**

A Horn clause is a clause with at most one positive literal. A positive Horn clause contains a positive literal, for example $[\neg P(x), P(f(x)]$ or $[P(a)]$. A negative Horn clause does not contain any positive literals, for example $[\neg P(x)]$ or $[\,]$.

(b) What is conflict resolution in production rule systems? State the conflict set and give an example of a conflict resolution strategy for the production system below: (7 marks)

Rule 1 $\forall x(InDoubt(x) \supset Scream(x))$

Rule 2 $\forall x(InDoubt(x) \supset Shout(x))$

Fact 1 $InDoubt(tom)$

Fact 2 $InDoubt(jerry)$

Answer. **K,C**

Several production rules may match the working memory, in several different ways, this generates several rule instances; conflict resolution is a procedure which is used to decide which instance to fire. (3 marks for explaining what conflict resolution is.) In the example, the conflict set is Rule 1 matched with x/tom, Rule 1 matched with $x/jerry$, Rule 2 matched with x/tom, Rule 2 matched with $x/jerry$. (2 marks for the conflict set.) An example of a conflict resolution strategy would be taking the first instance in lexicographic order (Rule 1 matched with x/tom), or stating which rule has higher priority. (2 marks for a concrete strategy.)

(c) Does forward chaining terminate if terms may contain function symbols? If yes, explain why, if not, give an example of a set of rules and facts for which forward chaining does not terminate. (5 marks)

Answer. **C**

In the presence of functions, forward chaining does not terminate; the simplest example is $\forall x(P(x) \supset P(f(x)))$ and $P(a)$. This generates an infinite sequence of facts $P(f(a))$, $P(f(f(a)))$, $P(f(f(f(a))))$, ...

[Note that this applies to the version of forward chaining as used in production systems, where new facts are added to the working memory. The propositional version from the textbook, where the facts are only marked as derived or not, does terminate. I gave answers based on that version 3 out of 5 marks. In the future, if using this question, I would make it more specifically apply to production systems.]

(d) This question is on propositional forward chaining. Given a set of n propositional rules of the form $p_1 \wedge \ldots \wedge p_k \supset q$ and m propositional facts of the form p_i, how many new facts can be generated (at most)? (3 marks)

Answer. **C**

Since there are n rules, and in propositional case each rule can match only one way, at most n rule instances will be fired all together. This means that at most n new facts can be generated.

(e) This question is on propositional resolution. Given a set of clauses which contain n different propositional symbols, is the search for the derivation of an empty clause by propositional resolution guaranteed to terminate, and if yes, after how many application of the resolution rule (at most)? (5 marks)

Answer. **K,C**

The search will terminate when we cannot derive any new clauses. The total number of clauses is the number of all possible sets of literals which can be made from n propositional variables. There are $2n$ literals (p and $\neg p$, for each propositional variable p). The number of different sets which contain these literals is 2^{2n}. So the search will terminate after 2^{2n} steps in the worst case.

4. This question is on ontologies and description logic

(a) What are ontologies? Give examples of widely used ontologies. (3 marks)

Answer. **K**

Ontologies are knowledge bases containing definitions, taxonomies and classifications. They are structure in terms of classes or concepts and relationships between them. Examples given in the course are Snomed, CYC, DBPedia, Google's Knowledge Graph.

(b) What is the relationship between W3C Web Ontology Language (OWL), OWL DL, and Description Logic? (5 marks)

Answer. **K,C**

OWL is based on RDF. OWL DL is a subset of OWL (with syntactic restrictions which make it decidable and semantically well defined) based on description logic ($SHOIN^D$: ALC with transitive roles, role hierarhies, nominals, inverse properties, cardinality restrictions, and data types).

(c) State the definition of concept construction in description logic ALC. Hint: assume an alphabet of atomic concepts $A_1, \ldots, A_n, \ldots$ and roles $r_1, \ldots, r_m, \ldots$ (5 marks)

Answer. **K**

- (sometimes used, although definable): $\top$, $\bot$ are concepts
- atomic concept is a concept
- if r is a role and C is a concept, then $\forall r.C$ is a concept (e.g. $\forall .child.Girl]$ describes someone all of whose children are girls)
- if r is a role and C is a concept, then $\exists r.C$ is a concept (e.g. $\exists .child.Girl]$ describes someone who has a daughter)
- if C is a concept then $\neg C$ is a concept
- if C_1 and C_2 are concepts then $C_1 \sqcap C_2$ is a concept
- if C_1 and C_2 are concepts then $C_1 \sqcup C_2$ is a concept

(d) Translate the following definitions and inclusion statements into ALC, using atomic concepts $Vegetarian$, $Person$, $Meat$, $Fish$, $Animal$, $Vegan$, and roles $Eats$ and $ProductOf$:

i. A vegetarian is a person who does not eat meat and does not eat fish (3 marks)

ii. A vegan is a person who does not eat animal products (3 marks)

iii. The concepts of non-vegetarian and vegan are disjoint (3 marks)

iv. Everyone who eats meat also eats animal products (3 marks)

Answer. **A**

i. $Vegetarian \equiv Person \sqcap \neg \exists Eats.Meat \sqcap \neg \exists Eats.Fish$

ii. $Vegan \equiv Person \sqcap \neg \exists Eats.\exists ProductOf.Animal$
($Vegan \equiv Person \sqcap \forall Eats.\forall ProductOf.\neg Animal$ is also fine)

iii. $\neg Vegetarian \sqcap Vegan \sqsubseteq \bot$ ($\neg Vegetarian \sqsubseteq \neg Vegan$ gets 2 marks)

iv. $\exists Eats.Meat \sqsubseteq \exists Eats.\exists ProductOf.Animal$