M2R Exam
Fundamentals of data processing and distributed knowledge
Semantics of Distributed Knowledge part

Duration : 3h
All documents allowed – *No* communication device allowed

January 2024

Note: Read all the questions carefully before answering. Justify your answers with respect to the semantics: this is the semantics that justifies their correctness.

Time and points are indicative.

Course questions [Expectation: 30mn; 5pts]

Here I give only three examples, but it should be around 10 questions, the answers are in the course in general.

Answers to these questions are generally short (if the answer is more than three sentences, it is probably wrong, except for the last one). They are related to the course content.

1. For what is it useful to query different sources?

   For accessing/retrieving more information.

2. What does it means for a structure (formula, theory, network, etc.) to be inconsistent?

   It has no model.

3. In modal logic, is $S \models S'$ defined by $\forall M, M \models S \Rightarrow M \models S'$ or by $\forall M, \forall w \in W_M, M, w \models S \Rightarrow M, w \models S'$? Does one expression implies the other? Why?

   The latter. It entails the former because it will also consider worlds in structures which do not universally satisfy S.

Application [Expectation: 2h; 15pts]

Data and ontology

We have three actors each one having their ontologies.

O_a: Cook

$$
\begin{aligned}
\text{oa:Meal} &\equiv \forall \text{oa:madeOf.oa:Ingredients} \\
\text{oa:VeggyMeal} &\equiv \text{oa:Meal} \sqcap \forall \text{oa:madeOf.}\neg\text{oa:Meat} \\
\text{oa:SpicyMeal} &\equiv \text{oa:Meal} \sqcap \exists \text{oa:madeOf.oa:Chilly} \\
\text{oa:Chily} &\sqsubseteq \text{oa:Ingredient} \\
\text{oa:Meat} &\sqsubseteq \text{oa:Ingredient}
\end{aligned}
$$

O_b: Producer

$$
\begin{aligned}
\top &\sqsubseteq \text{ob:Product} \\
\text{ob:Mushroom} &\equiv \text{ob:Product} \sqcap \text{ob:species} = \text{fungus} \\
\text{ob:Vegetable} &\equiv \text{ob:Product} \sqcap \text{ob:species} = \text{plant} \\
\text{ob:Meat} &\equiv \text{ob:Product} \sqcap \text{ob:species} = \text{animal} \\
\text{ob:Mushroom} &\perp \text{ob:Meat} \\
\text{ob:Vegetable} &\perp \text{ob:Meat} \\
\text{ob:RedMushroom} &\equiv \text{ob:Mushroom} \sqcap \text{ob:color} = \text{red} \\
\text{ob:WhiteMushroom} &\equiv \text{ob:Mushroom} \sqcap \text{ob:color} = \text{white} \\
\text{ob:BlackMushroom} &\equiv \text{ob:Mushroom} \sqcap \text{ob:color} = \text{black}
\end{aligned}
$$

O_d: Doctor / dietetician

$$
\begin{aligned}
\top &\sqsubseteq \text{od:Substance} \\
\text{od:Edible} &\sqsubseteq \text{od:Substance} \\
\text{od:Toxic} &\sqsubseteq \text{od:Substance} \\
\text{od:Lethal} &\sqsubseteq \text{od:Toxic} \\
\text{od:Edible} &\perp \text{od:Toxic}
\end{aligned}
$$

$\perp$ is owl:disjointWith, $\sqsubseteq$ is rdfs:subClassOf, $\equiv$ is rdfs:subClassOfowl:equivalentClass, $=$ is owl:hasValue, $\forall$ is owl:allValuesFrom, $\exists$ is owl:minQualifiedCardinality with $n = 1$ and $\sqcap$ is owl:intersectionOf.

This makes the following network:

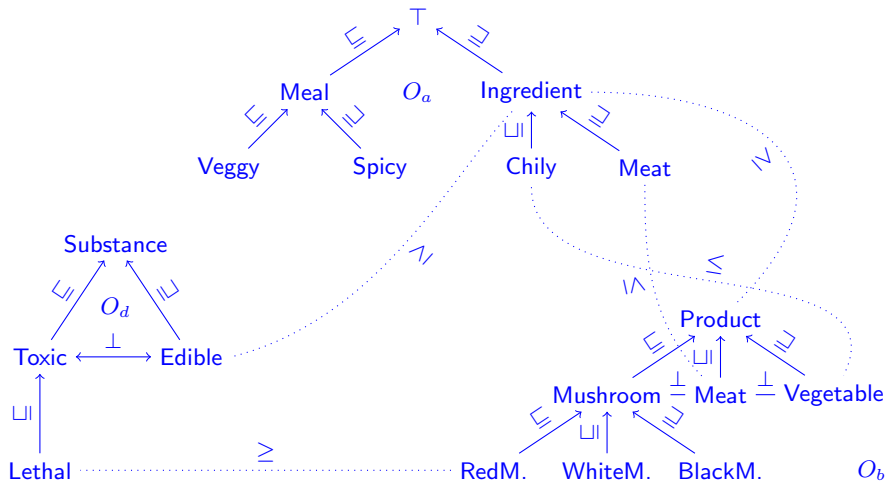

The following table features instances and their attribute values which make the graph G:

Instance in G	ob:taste	ob:color	ob:texture	ob:species
d:AmanitaP		white		fungus
d:AmanitaM		red		fungus
d:AjiA	spicy	yellow		plant
d:Snail		black	crisp	animal
d:Wasabi	spicy	green	crisp	
d:Shitake			velvety	fungus

4. Classify these instances, i.e. assign them to the classes in which they must belong in the ontologies above, i.e. for each $i \in G$, to all classes $c \in O$, such that $O \cup G \models_{OWL} i$ rdf:type c?

Instance in G	O_a	O_b	O_d
d:AmanitaP	$\top$	ob:WhiteMushroom	od:Substance
d:AmanitaM	$\top$	ob:RedMushroom	od:Substance
d:AjiA	$\top$	ob:Vegetable	od:Substance
d:Snail	$\top$	ob:Meat	od:Substance
d:Wasabi	$\top$	ob:Product	od:Substance
d:Shitake	$\top$	ob:Mushroom	od:Substance

5. Does $O_b \models_{OWL}$ ob:Vegetable$\bot$ob:Mushroom? Justify.

   No. Although, ob:Mushroom $\bot$ ob:Meat $\bot$ ob:Vegetable, the owl:disjointWith ($\bot$) relation is not transitive. Moreover, the owl:hasValue does not prevent to have several values for the same property. Hence, it is possible to have an individual related to both plant and fungus by the ob:species properties. Actually, it could also be possible that plant and fungus be interpreted as the same value and to build a model for these ontologies.

Alignments and queries

Consider the following data making G_a:

d:meal1 rdf:type oa:Meal
d:meal1 oa:madeOf d:AjiA
d:meal1 oa:madeOf d:Snail

d:meal2 rdf:type oa:VeggyMeal
d:meal2 oa:madeOf d:Wasabi
d:meal2 oa:madeOf d:Shitake

d:meal3 rdf:type oa:SpicyMeal
d:meal3 oa:madeOf d:Shitake
d:meal3 oa:madeOf d:AjiA

6. Does it allow you to reclassify some instances and which ones?

   Yes. d:AjiA, d:Snail, d:Wasabi and d:Shitake are now member of oa:Ingredient. Moreover, d:Wasabi and d:Shitake are now necessarily member of $\neg$oa:Meat.

7. Consider that a client wants a meal which is both spicy and vegeratian, how would you write such a query?

   SELECT ?x WHERE ?x rdf:type oa:SpicyMeal. ?x rdf:type oa:VeggyMeal.

8. Is this possible to find an answer to this query in one of the ontologies ($O_i \cup G_i$, $i \in \{a, b, d\}$ and G_i containing the statements using O_i vocabulary)?

   No. There is no individual related by rdf:type to both oa:SpicyMeal and oa:VeggyMeal.

Consider the following alignments between the ontologies:

$A_{a,b}$

oa:Meat $\leq$ ob:Meat
oa:Ingredient $\leq$ ob:Product
oa:Chily $\leq$ ob:Vegetable

$A_{b,d}$

ob:RedMushroom $\leq$ od:Lethal
ob:Meat $\leq$ od:Proteins

$A_{d,a}$

oa:Ingredient $\leq$ od:Edible

We consider the network of ontologies $\langle \Omega, \Lambda \rangle = \langle \{O_a \cup G_b, O_b \cup G_b, O_d\} \{A_{a,b}, A_{b,d}, A_{d,a}\} \rangle$.

9. Does it allow you to reclassify some instances with respect to $\langle \Omega, \Lambda \rangle$, i.e. $\langle \Omega, \Lambda \rangle \models i$ rdf:type c, and which ones?

   Yes. Now, all instances of oa:Ingredient, i.e. d:AjiA, d:Snail, d:Wasabi and d:Shitake, are now classified in od:Edible. Moreover, d:AmanitaM belongs to od:Lethal and hence od:Toxic.

10. Is there a query semantics under which your query of question 7 receives an answer, tell which answer(s) and explain why?

    No. In each of the meals, there is no way to know if we have all ingredients. And there is no way to know for the oa:VeggyMeal that some ingredients are oa:Chilly. Hence, even with the full network of ontologies, there is no way to return an answer to the query.

Revision and evolution

Now everyone becomes aware that Fred has eaten a delicate red mushroom. This means that the following statement has to be added to ontology O_d: d:AmanitaM rdf:type od:Edible.

11. Does this change the answer to the query of question 7? (please be precise: say in which interpretation, and how it would change the answer)

    If the whole network is used for interpreting the query. Since it is inconsistent, it entails any assignment of the variables as answer. At the very least, d:meal1, d:meal2 and d:meal3 are answers.

12. Is $O_d \cup \{$d:AmanitaM rdf:type od:Edible$\}$ inconsistent? Explain why.

    It is not inconsistent in itself. One can have a model in which this instance is interpreted as an od:Edible od:Substance.

13. Is the network with the assertion $(\langle\Omega, \Lambda\rangle \boxplus$ d:AmanitaM rdf:type od:Edible$/O_d)$ inconsistent? Explain why.

    The network is inconsistent, because d:AmanitaM would belong to the interpretation of od:Edible. But it belongs to the interpretation of ob:RedMushroom which by alignment $A_{b,d}$ belongs to od:Lethal with is included into od:Toxic which is disjoint from od:Edible.

14. Is there a local revision of ontology O_d that would restore consistency? If, yes provide a minimal one, if no, justify.

    Yes it is possible to suppress the assertion that Lethal $\sqsubseteq$ Toxic or that Edible$\perp$Toxic

15. Are there global revisions of $\langle\Omega, \Lambda\rangle$ that would restore consistency? If, yes provide a minimal one, if no, justify.

    Yes, in addition it would be possible to suppress the correspondence ob:RedMushroom $\leq$ od:Lethal from $A_{b,d}$ or the statements d:AmanitaM ob:color red, or d:AmanitaM ob:species fungus from O_b.

16. How to compute a partial meet revision of the network of ontologies by the assertion AmanitaM rdf:type od:Edible

    The point is to select some of the maximal subnetworks of $\langle\Omega, \Lambda\rangle$ consistent with d:AmanitaM rdf:type od:Edible. The maximal consistent subnetworks are all those networks $\langle\Omega, \Lambda\rangle$ to which one of the statement or correspondence in the answers to questions 14 and 15 has been suppressed. One basic idea would be to select a local revision if there exist one. But as seen above, this is not really satisfying. Then an alternative solution would be to modify the alignments as they are usually less "correct" than the ontologies. Here it is clear that the incorrect statement is ob:RedMushroom $\leq$ od:Lethal (this is indeed not true) from $A_{b,d}$, but there is no real way to select it.

Open question [Expectation: 15mn; 3pts]

Multi-agent epistemic and doxastic logics express what agents know and believe. It should be possible to turn a network of ontology into a theory in such a logic. The logic comes with a semantics defining how the theory can entail ($\models_{MAEL}$) statements.

How would you use these considerations to define the semantics of query evaluation by using such techniques? What problem does it pose? (No need to be formal for this question)

Usually, the answers to a query $q[\vec{x}]$ are those assignments σ of the variables of $\vec{x}$ that make $\sigma(q[\vec{x}])$ entailed by the network. Hence, a possible solution would be to use the entailment of the logic ($\models_{MAEL}$). But this requires that both the network of ontologies and the query be transformed into the logic. One may use the transformation τ provided in the last part of the course. However, it distinguishes between knowledge (the ontologies) and beliefs (the alignments), which is not necessary here. Since we want to evaluate the query assuming the truthfulness of ontologies and alignments, it is possible to transform every statement into knowledge.