G53KRR SEM Feedback and Revision 2017

SEM feedback

- No big issues raised
- I will go through SEM questions one by one

Q1

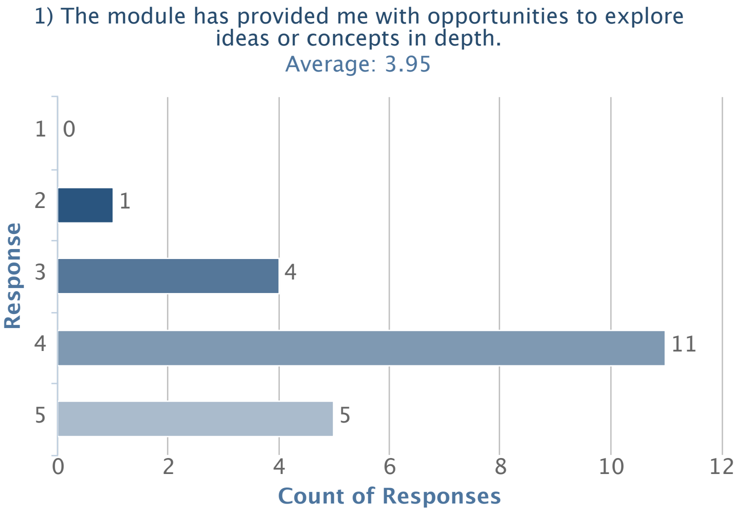
1) The module has provided me with opportunities to explore ideas or concepts in depth.
Average: 3.95
Response
1
2
3
4
5
0
1
4
11
5
0
2
4
6
8
10
12
Count of Responses

Q1

- I hope that you had an opportunity to study formal languages for knowledge representation in depth
- the fundamental concept of logical entailment
- the idea of logical reasoning/formal derivation

Q2

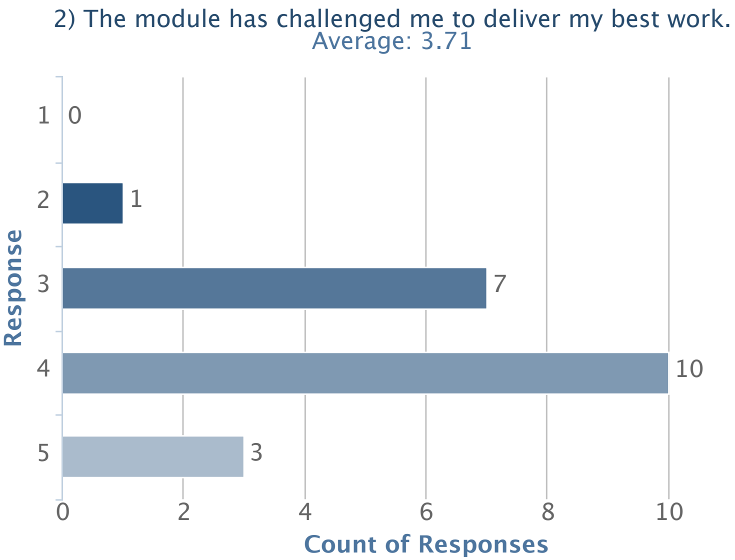
2) The module has challenged me to deliver my best work.
Average: 3.71
Response
1
2
3
4
5
0
1
7
10
3
0
2
4
6
8
10
Count of Responses

Q2

- this is not easy to answer until the exam...

Q3

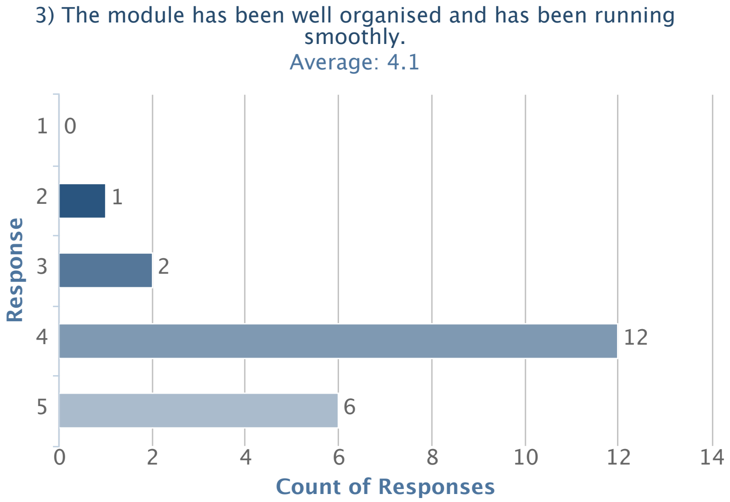

3) The module has been well organised and has been running smoothly.
Average: 4.1
Response
1
2
3
4
5
0
1
2
12
6
0
2
4
6
8
10
12
14
Count of Responses

Q3

- did not get a text comment what was bad about organisation of the module
- would be grateful for feedback, anonymous or not

Q4

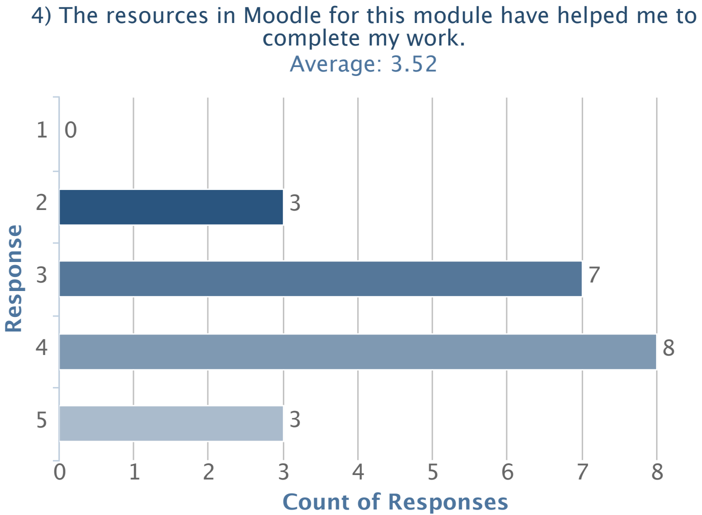
4) The resources in Moodle for this module have helped me to complete my work.
Average: 3.52
Response
Count of Responses
1
2
3
4
5
0
3
7
8
3
0 1 2 3 4 5 6 7 8

Q4

- there were clearly no resources on moodle apart from Logic in CS book
- there are however resources on the module web page, www.cs.nott.ac.uk/~psznza/G53KRR which is linked from moodle
- hope everyone found it – it is essential for revision

Q5

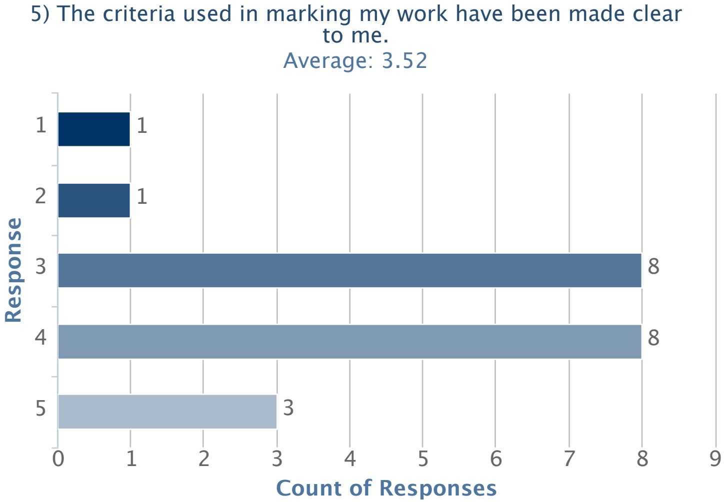
5) The criteria used in marking my work have been made clear to me.
Average: 3.52
Response
1
2
3
4
5
1
1
8
8
3
0
1
2
3
4
5
6
7
8
9
Count of Responses

Q5

- exam papers are on the web
- sample model answers and marking schemes are on the web
- any correct answer will get full marks
- small mistakes will get very small mark reductions or none at all, if there is evidence of clear understanding
- if you want to do a sample exam/informal coursework and get my feedback, please email it to me

Q6

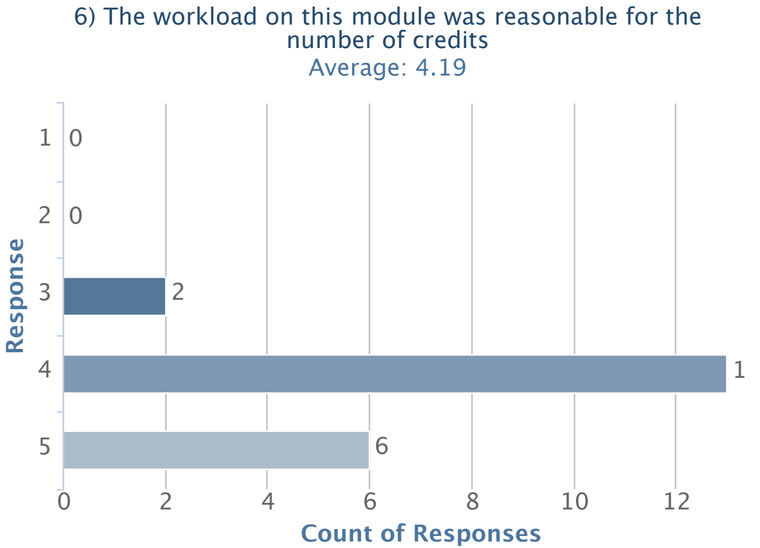
6) The workload on this module was reasonable for the number of credits
Average: 4.19
Response
1
2
3
4
5
0
0
2
1
6
0
2
4
6
8
10
12
Count of Responses

Q6b

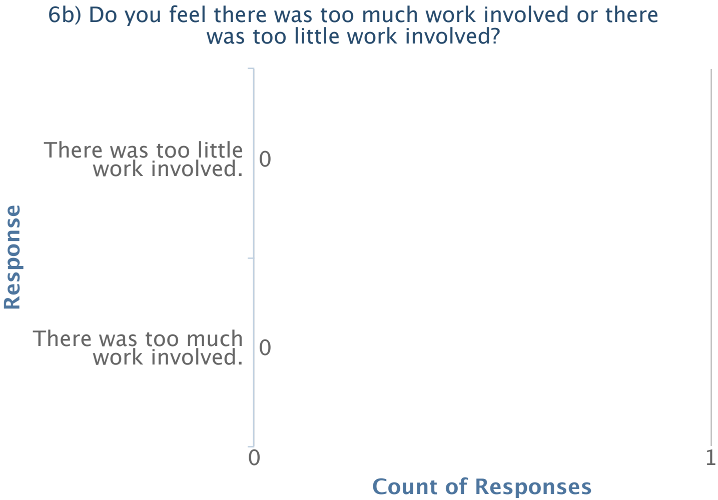
6b) Do you feel there was too much work involved or there was too little work involved?
There was too little work involved.
0
There was too much work involved.
0
Response
0
1
Count of Responses

Q7

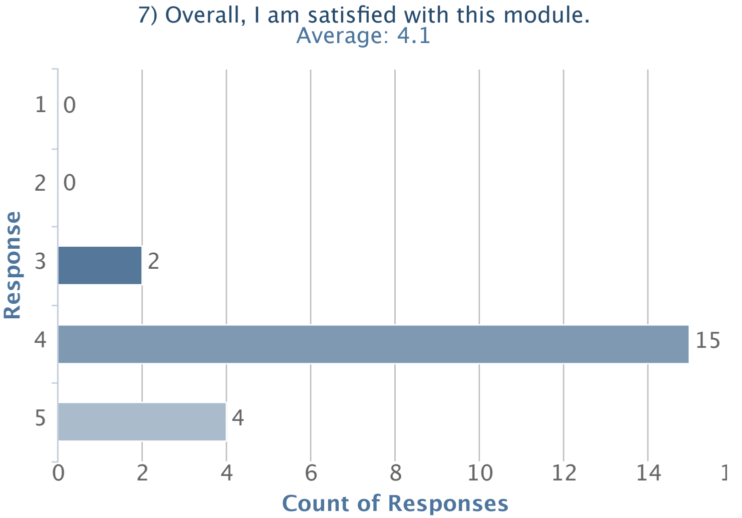

7) Overall, I am satisfied with this module.
Average: 4.1
Response
1
2
3
4
5
0
0
2
15
4
0
2
4
6
8
10
12
14
1
Count of Responses

Exam format

- same as in 2014 - 2016: 4 questions out of 4
- previous papers and answers on the G53KRR module web page

Exam marking

- Answers to the exams published on the web are sample answers
- Often there are many different correct answers possible (for example, resoulution derivations)
- Any correct answer will get full marks
- Answers which have minor mistakes will have minor marks reductions (or sometimes no reduction, if the answer shows clear understanding)
- Big mistakes which suggest poor understanding will lose more marks

Exam topics = overview of the module

- First order logic
- Resolution
- Horn clauses, backward chaining, forward chaining
- Knowledge elicitation, lifecycle of a knolwedge based system
- Description logic
- Non-monotonic reasoning
- Bayesian networks

How to revise

- For each lecture, check recommended reading (on G53KRR web page)
- Do informal exercises (on G53KRR web page)
- Do exam questions from previous years (on G53KRR web page)
- Revision lecture in January?
- Send me topics for revision in advance
- Send me your exam attempts for feedback

Some common mistakes

- Next I will go through some common mistakes, topic by topic

Common mistakes 1

First order logic question: show that S1, S2 do not logically entail S3

- Correct answer: describe an interpretation which makes S1 and S2 true and S3 false
- Don't:
  - use truth tables for first order sentences
  - attempt a resolution derivation of [] from S1, S2 and ¬S3, and then stop and say 'see, it does not work, so S3 is not entailed'

Common mistakes 2

Resolution

- don't apply resolution to two literals at the same time:

$$MISTAKE: \frac{[A, B],\ [\neg A, \neg B]}{[\,]}$$

it is not sound! $A \vee B$ and $\neg A \vee \neg B$ should not derive false.

- only substitute for variables (not constants or functional terms) Don't do $f(x)/a$ or $a/f(x)$.

$$MISTAKE: \frac{[(P(f(x))],\ [\neg P(a)]}{[\,]}$$

Bayesian networks

- Directed acyclic graph
- Nodes: propositional variables; a directed edge from p_i to p_j if the truth of p_i affects the truth of p_j. p_i parent of p_j.
- $$J(\langle P_1, \ldots, P_n \rangle) = Pr(P_1 \wedge \ldots \wedge P_n)$$
- Chain rule
$$Pr(P_1 \wedge \ldots \wedge P_n) = Pr(P_1) \cdot Pr(P_2 | P_1) \cdots Pr(P_n | P_1 \wedge \ldots \wedge P_{n-1})$$
- Independence assumption *Each propositional variable in the belief network is conditionally independent from non-parent variables given its parent variables:*
$$Pr(P_i \mid P_1 \wedge \ldots \wedge P_{i-1}) = Pr(P_i \mid parents(P_i))$$
where $parents(P_i)$ is the conjunction of literals which correspond to parents of p_i in the network.

Mistake 3

- Mistake: suppose a network consists of two variables, p_1 and p_2, such that there is an edge from p_1 to p_2. The mistake is to say that $Pr(p_1 \mid p_2) = Pr(p_1)$ because p_2 is not a parent of p_1 (so apply the independence assumption 'in reverse order of indices')
- This is a much more subtle (and understandable given the way the independenc assumption is stated) mistake.
- The independence assumption statement assumes that in the state description, the variables are listed in topological sort order (if there is an edge from p_i to p_j, then p_i appears before p_j in the order). This is always possible since the graph is acyclic. So we never check probability of parent conditioned on a child or a set of descendants.

Some revision slides

- Next I will go through some revision topics and previous exam answers

Circumscription

- The main idea: formalise common sense rules which admit exceptions.
- Rules like 'Birds fly' formalised as

$$\forall x(Bird(x) \wedge \neg Ab(x) \supset Flies(x))$$

- To check whether something is entailed by a knowledge base which contains such rules, we only check if it is entailed under the assumption that the set of exceptions $I(Ab)$ is as small as possible
- This is called circumscription or minimal entailment

Example

$$KB = \{Bird(tweety), \forall x(Bird(x) \wedge \neg Ab(x) \supset Flies(x))\}$$

- Classically, $KB \not\models Flies(tweety)$ because there are interpretations of *KB* where Tweety is an exceptional bird (it is in $I(Ab)$) and it does not fly
- But such interpretations do not minimise the set of exceptions: nothing which is said in *KB* forces us to think that Tweety is exceptional, so it does not have to be in $I(Ab)$
- If we only consider interpretation which satisfy *KB* and where the set of exceptions is as small as possible, Tweety is not in this set, so $Bird(tweety) \wedge \neg Ab(tweety)$ holds and hence $Flies(tweety)$ holds
- *KB* entails $Flies(tweety)$ on 'minimal' interpretations where $I(Ab)$ is circumscribed (made as small as possible)

Definition of minimal entailment

- Let $M_1 = (D, I_1)$ and $M_2 = (D, I_2)$ be two interpretations over the same domain such that every constant and function are interpreted the same way.
- $M_1 \leq M_2 \ \Leftrightarrow \ I_1(Ab) \subseteq I_2(Ab)$
- $M_1 < M_2$ if $M_1 \leq M_2$ but not $M_2 \leq M_1$. (There are strictly fewer abnormal things in M_1).
- *Minimal entailment:* $KB \models_{\leq} \alpha$ iff for all interpretations M which make KB true, either $M \models \alpha$ or M is not minimal (exists M' such that $M' < M$ and $M' \models KB$).

Back to the example

- $KB = \{Bird(tweety), \forall x(Bird(x) \wedge \neg Ab(x) \supset Flies(x))\}$
- $KB \models_{\leq} Flies(tweety)$ because for every interpretation M which makes KB true and $Flies(tweety)$ false, it has to be that $I(tweety) \in I(Ab)$.
- So for for every such interpretation there is an interpretation M' which is just like M, but $I'(tweety) \notin I'(Ab)$ and $I'(tweety) \in I'(Flies)$, and M' still makes KB true and $M' < M$.

Defaults

- A *default rule* consists of a *prerequisite* α, *justification* β, *conclusion* γ and says 'if α holds and it is consistent to believe β, then believe γ': $\frac{\alpha : \beta}{\gamma}$
- For example:

$$\frac{Bird(x) : Flies(x)}{Flies(x)}$$

- Default rules where justification and conclusion are the same are called *normal default rules* and are writted $Bird(x) \Rightarrow Flies(x)$.

Default theories and extensions

- A default theory *KB* consists of a normal first-order knowledge base *F* and a set of default rules *D*
- A set of reasonable beliefs given a default theory $KB = \{F, D\}$ is called an *extension* of *KB*
- *E* is an *extension* of (F, D) iff for every sentence π,

$$\pi \in E \;\Leftrightarrow\; F \cup \{\gamma \mid \frac{\alpha : \beta}{\gamma} \in D, \alpha \in E, \neg\beta \notin E\} \models \pi$$

How one could construct an extension

$$\pi \in E \ \Leftrightarrow\ F \cup \{\gamma \mid \frac{\alpha : \beta}{\gamma} \in D, \alpha \in E, \neg\beta \notin E\} \models \pi$$

1. $E := F$;
2. close E under classical entailment: $E := \{\pi : E \models \pi\}$
3. choose some (substitution instance of) $\frac{\alpha:\beta}{\gamma} \in D$
4. if $\alpha \in E$, and $\neg\beta \notin E$ (meaning, β is consistent with E), $E := E \cup \{\gamma\}$
5. go back to 2

Example

$$F = \{Bird(tweety)\}, D = \{\frac{Bird(x) : Flies(x)}{Flies(x)}\}$$

- $E := \{Bird(tweety)\}$
- close E under classical entailment: $E := \{\pi : Bird(tweety) \models \pi\}$
- $\frac{Bird(tweety):Flies(tweety)}{Flies(tweety)} \in D$
- $Bird(tweety) \in E$, and $\neg Flies(tweety) \notin E$
  $E := E \cup \{Flies(tweety)\}$
- $E := \{\pi : Bird(tweety), Flies(tweety) \models \pi\}$
- there are no more rules to apply

Example from 2008 exam, Q6e

$$F = \{Dutchman(peter),\ Dutchman(hans),\ Dutchman(johan),$$

$$peter \neq hans,\ hans \neq johan,\ peter \neq johan,$$

$$\neg Tall(peter) \vee \neg Tall(hans)\}$$

$$D = \{\frac{Dutchman(x) : Tall(x)}{Tall(x)}\}$$

Three instances of the default rule:

$$\frac{Dutchman(peter) : Tall(peter)}{Tall(peter)} \quad \frac{Dutchman(hans) : Tall(hans)}{Tall(hans)}$$

$$\frac{Dutchman(johan) : Tall(johan)}{Tall(johan)}$$

Exam 2008 example continued

- Suppose we start constructing E_1 with the first rule, for Peter. Since $\neg Tall(peter) \notin E_1$, we can add $Tall(peter)$ to E_1.
- After we close E_1 under consequence, from $Tall(peter)$ and $\neg Tall(peter) \vee \neg Tall(hans)$ we get $\neg Tall(hans) \in E_1$.
- So now the second rule for Hans is not applicable.
- The third rule is applicable, since $\neg Tall(johan) \notin E_1$, we can add $Tall(johan)$ to E_1
- Another possible extension is E_2: we use the second rule first, and add $Tall(hans)$ to E_2.
- Now the first rule is not applicable, because E_2 contains $\neg Tall(peter)$
- The third rule is applicable, since $\neg Tall(johan) \notin E_2$, we can add $Tall(johan)$ to E_2

Another example with two extensions

- Facts: $F = \{Republican(dick), Quaker(dick)\}$
- Default rules: $Republican(x) \Rightarrow \neg Pacifist(x)$, $Quaker(x) \Rightarrow Pacifist(x)$.
- Extension E_1 (pick the rule $Republican(x) \Rightarrow \neg Pacifist(x)$ first) is all consequences of $\{Republican(dick), Quaker(dick), \neg Pacifist(dick)\}$. Because we start with $E_1 = \{Republican(dick), Quaker(dick)\}$, $\neg\neg Pacifist(dick) \not\in E_1$, so we can add $\neg Pacifist(dick)$ to E_1.
- Extension E_2 (pick the rule $Quaker(x) \Rightarrow Pacifist(x)$ first) is all consequences of $\{Republican(dick), Quaker(dick), Pacifist(dick)\}$. Because we start with $E_2 = \{Republican(dick), Quaker(dick)\}$, $\neg Pacifist(dick) \not\in E_2$, so we can add $Pacifist(dick)$ to E_2.

Example with one extension

- Facts: $F = \{Republican(dick), Quaker(dick), \forall x(Republican(x) \supset MemberOfPoliticalParty(x))\}$
- Default rules: $Republican(x) \Rightarrow \neg Pacifist(x)$,

$$\frac{Quaker(x) : Pacifist(x) \wedge \neg MemberOfPoliticalParty(x)}{Pacifist(x)}$$

- Closure of F under consequence includes: $\{Republican(dick), Quaker(dick), \forall x(Republican(x) \supset MemberOfPoliticalParty(x)), MemberOfPoliticalParty(dick)\}$
- The second default rule is not applicable, because $\neg\neg MemberOfPoliticalParty(dick) \in E$
- only the first rule is applicable, since $\neg\neg Pacifist(dick) \notin E$, so $\neg Pacifist(dick)$ is added.

Any questions?