AI Techniques A Exercise Sheet 8 – Marking Guidelines

Answers for 90 minutes under closed book conditions

1. (a) Some proponents of AI like to distinguish between scientific and engineering goals. Explain that distinction and comment briefly on whether you think it is useful. [3%]

<u>Scientific Goal</u> : To determine which ideas about representing knowledge, search, learning, rule use, and so on, explain various sorts of real intelligence, e.g. how real brains work.

<u>Engineering Goal</u> : To solve real world problems using AI techniques such as representing knowledge, search, learning, rule use, and so on.

The distinction can be <u>useful</u>, since there are clearly different objectives, but it can also be <u>not useful</u>, because common techniques are involved and there can be advantages to having a flow of ideas and results between the two sub-fields.

(b) The field of AI has its roots in several older disciplines. List the principal ones and outline one important idea that each brings to the study of AI. [7%]

The main underlying disciplines are: Philosophy, Logic/Computation, Psychology/ Cognitive Science, Biology/Neuroscience, Evolution. Listing four non-overlapping ones, and giving any sensible "important ideas" (e.g. concerning computability, inference, representation, learning, neural networks, …) will get the marks.

2. (a) What is an "agent"? What is a "rational agent"? [2%]

From Russell & Norvig, page 31: "An *agent* is anything that can be viewed as perceiving its environment through *sensors* and *acting* upon that environment through *effectors*".

A *rational agent* is one that acts in a manner that causes it to be the most successful as measured by an appropriate performance measure.

(b) In describing intelligent agents it is often convenient to specify them in terms of Percepts, Actions, Goals and Environment. State briefly what each of these four concepts mean. [4%]

Percepts – the inputs that the agent receives, e.g. through its senses

Actions – the outputs that the agent produces

Goals – what the agent is expected/aiming to achieve

Environment – the external system(s) that the agent is interacting with

(c) List what these concepts correspond to in the following agents:

(i) A medical diagnostic system.

(ii) An object sorting robot. [4%]

Agent Type	Percepts	Actions	Goals	Environment
Medical diagnostic system	Symptoms, test results, patient's answers	Questions, test requests, treatments	Healthy patients, minimise costs	Patient, hospital
Object sorting robot	Pixels of varying intensity and colour	Pick up parts and sort into bins	Place parts into correct bin	Conveyor belt with parts

3. (a) Represent the following knowledge in a semantic network:

Dogs are Mammals	**Birds are Bipeds**
Mammals are Animals	**Humans are Bipeds**
Birds are Animals	**Humans are Mammals**
Fish are Animals	**Dogs chase Cats**
Worms are Animals	**Cats eat Fish**
Cats are Mammals	**Birds eat Worms**
Cats are Quadrupeds	**Fish eat Worms**
Dogs are Quadrupeds	

[6%]

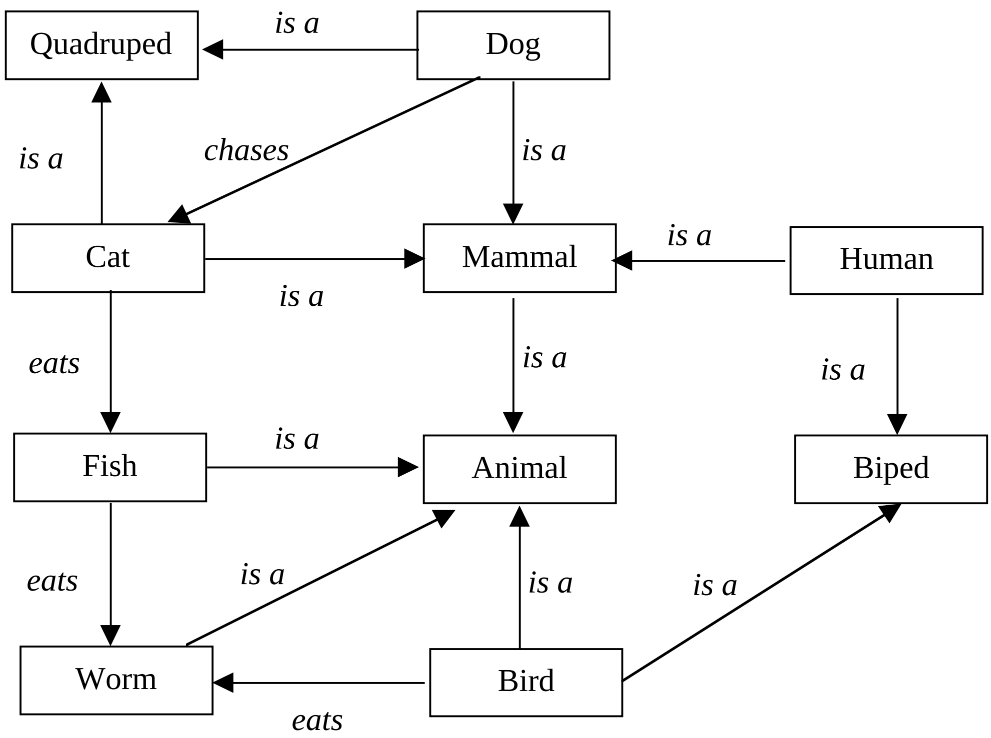

(b) **Suppose you learn that *Lassie* is a dog. What additional knowledge about *Lassie* can be derived from your representation? Explain how.** **[2%]**

Lassie inherits all the properties of dogs from further up the hierarchy of links. So Lassie is a quadruped, is a mammal, is an animal, and chases cats. Zero marks if Lassie is said to eat fish.

(c) **What is a tangled hierarchy? Why can they result in conflicts? How can these conflicts be resolved? Illustrate your answers with a simple example.** **[7%]**

A tangled hierarchy is a sequence of semantic net connections that is not a simple tree. They can result in conflicts if different/inconsistent defaults can be inherited by a single instance. For example, consider the question: "Can Oli the ostrich fly?"

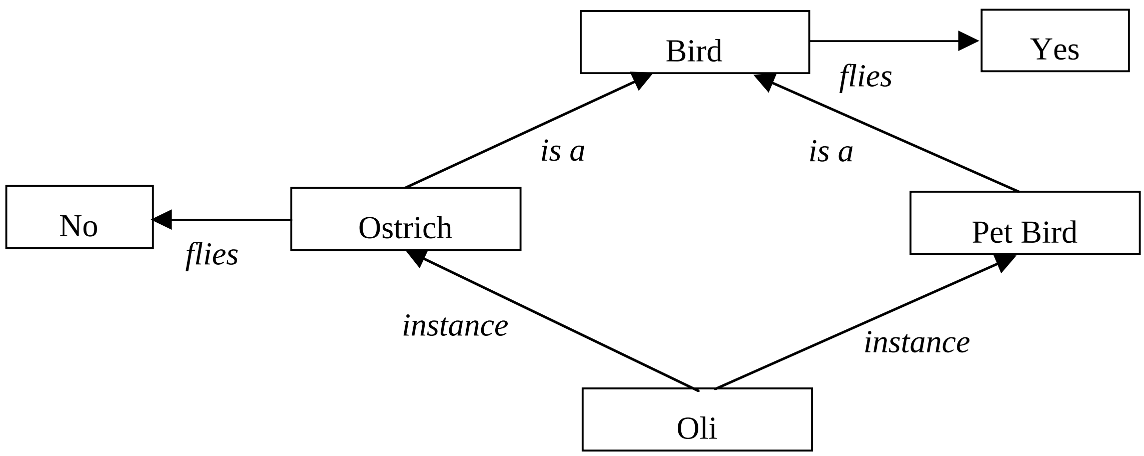

The left hand branch says Oli cannot fly, the right hand branch says he can.

One solution to the inheritance conflict would be to attach a specific "flies no" for all individual instances of an ostrich. A much better solution would be to have an algorithm for traversing the hierarchy which guarantees that specific knowledge will always dominate over general knowledge. A concept of "inferential distance" which defines closeness will allow the inferential engine to resolve the conflict appropriately.

4. The following production system was designed to help a new zoo-keeper look after his animals. The notation used is such that "*x*" stands for an animal, "bird(*x*)" means "*x* is a bird", and so forth.

Rules:

R1:

IF:	feathers(*x*)
THEN:	bird(*x*)

R2:

IF:	flies(*x*) & lays_eggs(*x*)
THEN:	bird(*x*)

R3:

IF:	gives_milk(*x*)
THEN:	mammal(*x*)

R4:

IF:	eats_meat(x)
THEN:	carnivore(*x*)

R5:

IF:	mammal(*x*) & sharp_teeth(*x*)
THEN:	carnivore(*x*)

R6:

IF:	carnivore(*x*)
THEN:	feed_meat(*x*)

R7:

IF:	bird(*x*) & not_flies(*x*)
THEN:	penguin(*x*)

R8:

IF:	penguin(*x*)
THEN:	feed_fish(*x*)

R9:

IF:	carnivore(*x*)
THEN:	dangerous(*x*)

Initial facts:

sharp_teeth(Lucy)	feathers(Penny)	not_flies(Penny)
gives_milk(Lucy)	lays_eggs(Penny)	

(a) What is "binding"? How is the "conflict set" defined in general? What is the initial conflict set in the above example? [3%]

Binding is the process of associating values to variables in a rule, e.g. setting x = Lucy in R3.

The conflict set is the set of rules which can fire given the facts in working memory. These will be a set of pairs of the form:

< Production rule ; Matching facts in working memory >

The initial conflict set is the set of rules that can be fired with the initial facts. In this case, it is:

< R1 ; feathers(Penny) >
< R3 ; gives_milk(Lucy) >

(b) How should you resolve the conflict in this case? Give reasons for your answer. [2%]

Given that all the rules in the conflict have the same age, have not been used before, and have equal specificity (number of pre-conditions), and we have no goal, random choice is a reasonable way to resolve the conflict. However, assuming you care for the safety of the zoo-keeper, you should notice that R3 might lead to dangerous(x) whereas R1 will not, so firing R3 first would be a sensible idea.

(c) What can be derived from the knowledge base by forward reasoning? Explain your answer in detail. [5%]

Firing both rules R1 and R3 gives new facts:

R1 => bird(Penny)
R3 => mammal(Lucy)

This allows two more rules to fire:

< R5 ; mammal(Lucy) , sharp_teeth(Lucy) >
< R7 ; bird(Penny) , not_flies(Penny) >

which lead to new facts:

R5 => carnivore(Lucy)
R7 => penguin(Penny) ,

This allows two more rules to fire:

< R6 ; carnivore(Lucy) >
< R8 ; penguin(Penny) >

< R9 ; carnivore(Lucy) >

which lead to new facts:

R6 => feed_meat(Lucy)
R8 => feed_fish(Penny)
R9 => dangerous(Lucy)

These seven new facts are what can be derived from the knowledge base.

(d) How can backward reasoning be used to determine which animals are known to be dangerous? Work through the details for the above case. [5%]

Our goal is dangerous(x). Backward chaining works by repeatedly looking for goals/sub-goals in the THEN parts of the rules, and taking the corresponding IF parts of the rules as new sub-goals. When there are no more sub-goals, we insert the initial facts to determine which sub-goals/goals can fire.

To find the dangerous animals we start with the only rule involving "dangerous":

R9 dangerous(x) IF carnivore(x)

Then two rules imply "carnivore":

R4 carnivore(x) IF eats_meat(x)
R5 carnivore(x) IF mammal(x) AND sharp_teeth(x)

Then only one rule implies any of "eats_meat", "mammal", or "sharp teeth":

R3 mammal(x) IF gives_milk(x)

No further rules are relevant. We now reverse the process inserting our initial and derived known facts:

R3 gives_milk(Lucy) => mammal(Lucy)
R5 mammal(Lucy) AND sharp_teeth(Lucy) => carnivore(Lucy)
R4 can't fire
R9 carnivore(Lucy) => dangerous(Lucy)

We conclude that Lucy is known to dangerous. We do not conclude that Penny is known to be dangerous, but that does not necessarily mean that Penny is not dangerous.