CS 344 (Spring 2018): Class Test 1

Instructor: Shivaram Kalyanakrishnan

11.05 a.m. – 12.00 p.m., January 31, 2018, 103 New CSE Building

Total marks: 15

Note. Provide brief justifications and/or calculations along with each answer to illustrate how you arrived at the answer.

Question 1.

1a. Describe the Turing Test and explain its objective. [2 marks]

1b. The RoboCup Federation has set itself the target that

> "by the middle of the 21st century, a team of fully autonomous humanoid robot soccer players shall win a soccer game, complying with the official rules of FIFA, against the winner of the most recent World Cup."

What similarities and differences do you observe between the Turing Test and the RoboCup challenge? [2 marks]

Question 2. Consider two search problem instances, X_1 and X_2, which are identical except for their cost functions, which are c_1 and c_2, respectively. Heuristic function h_1 is consistent with respect to c_1, and heuristic function h_2 is consistent with respect to c_2.

2a. Consider a third search problem instance X_3, which is identical to X_1 and X_2 except for its cost function c_3, given by $c_3 = \max\{c_1, c_2\}$. Is the heuristic function $h_3 = \max\{h_1, h_2\}$ guaranteed to be consistent with respect to c_3? Prove that your answer is correct. [2 marks]

2b. Consider a fourth search problem instance X_4, which is identical to X_1 and X_2 except for its cost function c_4, given by $c_4 = \min\{c_1, c_2\}$. Is the heuristic function $h_4 = \min\{h_1, h_2\}$ guaranteed to be consistent with respect to c_4? Prove that your answer is correct. [2 marks]

Question 3. The following questions pertain to a turn-taking game played by three agents: A, B, and C. The corresponding game tree is shown on the next page. Each square corresponds to a node, and shows the agent in charge of making a move at that node. The actions available at each such node are "left" (L) and "right" (R). Chance nodes are shown as circles; edges emanating from them correspond to random events (denoted $E1$, $E2$, $E3$, and $E4$), whose probabilities are specified alongside. Each leaf contains a triple (x, y, z), indicating that if the leaf is reached, A gets a reward of x, B gets a reward of y, and C gets a reward of z.

Observe that A makes the first move. B and C do not necessarily get to make moves in all the games. Regardless, all three players are given their rewards when a leaf node is reached.

For your convenience, you are given two copies of the same figure; feel free to scribble on them for your rough work. While writing down your solutions, name your nodes by the path leading to them from the root. For example, the root node is $\emptyset$, and the left-most leaf is L-E1-L-L.

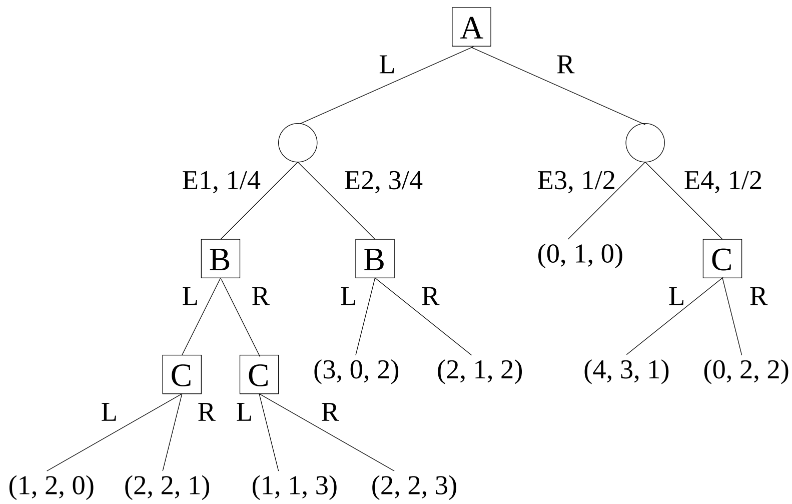

3a. If A plays action L from the root node, while B and C both pick actions uniformly at random from all nodes, what is B's expected reward? [2 marks]

3b. Suppose A and C collude against B: that is, whatever B's strategy, A and C will take actions to minimise B's expected reward from that strategy. If so, what strategy must B follow to maximise its own expected reward? What are A's and C's strategies? What is the resulting expected reward for B? (To specify a player's *strategy*, you have to provide an action for each of the player's decision-making nodes.) [3 marks]

3c. Suppose the three players decide to *cooperate*, and take actions such that the expected sum of their rewards is maximised. Which action is taken at each node? What is the expected sum of rewards obtained? [2 marks]

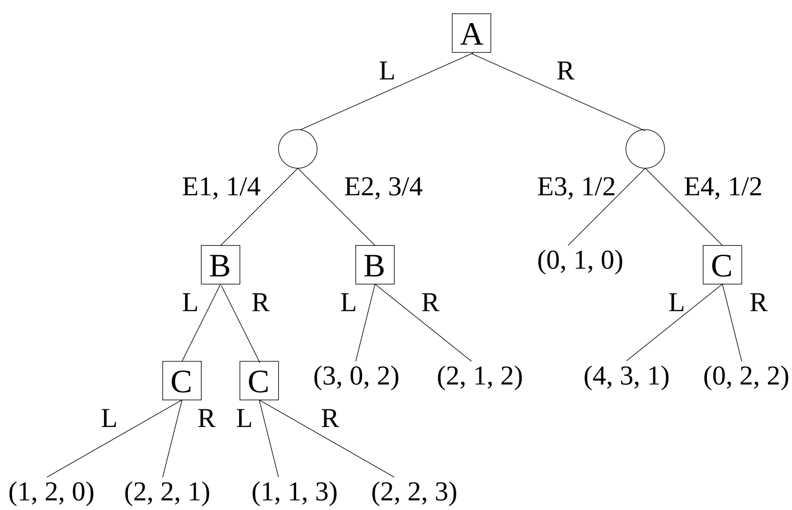

Solutions

1a. The Turing Test was originally proposed by Alan Turing as a test of the proficiency achieved by the field of AI. An AI agent A passes the test if a human judge J is unable to distinguish it from a human agent H. To date no AI agent has passed the test.

In the test, A and H are kept in separate rooms. J sits in another separate room, which is connected with the rooms of A and H using text-based communication lines. J does not know which line connects to which agent—it is precisely this information that it must attempt to glean. In order to do so, J can send questions on each line, and examine the answers returned. J is free to ask any questions in order to arrive at its conclusion. If J arrives at the correct conclusion, A fails the test; else A passes.

1b. To pass the Turing Test, an agent does not have to take physical actions with its body: it solely has to send signals based on its thinking. Soccer, however is an intrinsically *embodied* enterprise. A "brain" alone will not suffice: in order to succeed at soccer, the agent has to execute actions (such as moving and kicking a ball) in the real world, subjected to the laws of physics.

The two tests have in common that they define a benchmark for AI based on a comparison with *human* abilities. The need for embodiment implies that the RoboCup agent will be resource-constrained: it cannot use a large computing server, while agent A in the Turing Test agent. Another distinguishing attribute of the RoboCup challenge is that it involves multiple AI agents.

2a. The heuristic h_3 is guaranteed to be consistent, since for every n, a, n', we have

$$\begin{aligned} h_3(n) &= \max\{h_1(n), h_2(n)\} \\ &\leq \max\{c_1(n, a, n') + h_1(n'), c_2(n, a, n') + h_2(n')\} \\ &\leq \max\{c_1(n, a, n'), c_2(n, a, n')\} + \max\{h_1(n'), h_2(n')\} \\ &= c_3(n, a, n') + h_3(n'). \end{aligned}$$

2b. The heuristic h_4 need not be consistent. The search instance shown below, along with particular choices of c_1, h_1 and c_2, h_2 serves as a counterexample. Observe that h_1 and h_2 are consistent with respect to c_1 and c_2, respectively, but $h_4 = \min\{h_1, h_2\}$ is not consistent with respect to $c_4 = \min\{c_1, c_2\}$.

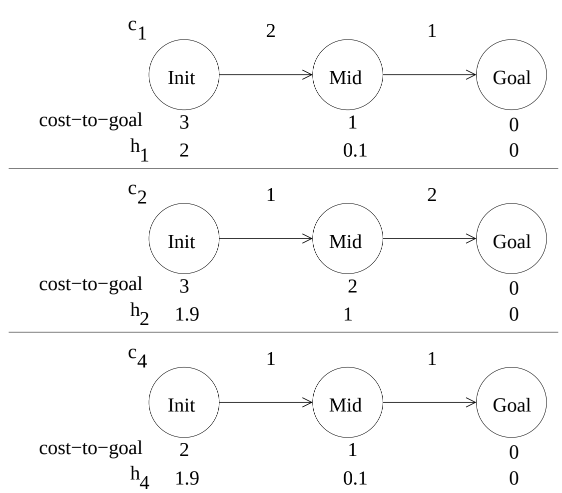

3a. Since A plays L, only the left sub-tree can be accessed. The table below shows the probability of reaching each leaf in the left sub-tree and the associated reward for B.

Node	Probability	B's reward
L-E1-L-L	1/16	2
L-E1-L-R	1/16	2
L-E1-R-L	1/16	1
L-E1-R-R	1/16	2
L-E2-L	3/8	0
L-E2-R	3/8	1

Weighting the rewards by the probabilities and summing, we find that B's expected reward is $\frac{13}{16}$.

3b. B has no role to play in the right sub-tree. If A plays R, and C plays R from R-E4, they can keep B's expected reward down to $\frac{1}{2}(1) + \frac{1}{2}(\min\{3, 2\}) = 1.5$. The only reason for A to play L would be if B can be forced to an even lower expected reward in the left sub-tree. Indeed it can.

By playing adversarially, C can restrict B's reward to 2 from L-E1-L, and to 1 from L-E1-R. Thus, the highest B can achieve from L-E1 is 2. Similarly, the highest B can achieve from L-E2 is $\max\{0, 1\} = 1$. Hence, if A plays L and C cooperates with A to minimise B's reward, the best B can achieve is $\frac{1}{4}(2) + \frac{3}{4}(1) = 1.25$.

Therefore: A plays L; C plays either L or R from L-E1-L since they give the same reward to B; and C plays L from L-E1-R. The right sub-tree is not visited (and so A's and C's actions in this portion need not be specified). If A and C play in this manner, B can at most achieve an expected reward of 1.25—it achieves this reward by playing L from L-E1 and R from L-E2. If B plays any other strategy, A and C can contain it to less than 1.25.

3c. If a particular node is reached, it is clear that the outgoing action must be the one that will yield the highest expected sum of rewards. It is easy to fix this choice at the nodes that only lead to leaves:

- from L-E1-L, the optimal action is R, which gives a combined reward of $2 + 2 + 1 = 5$;
- from L-E1-R, the optimal action is R, which gives a combined reward of $2 + 2 + 3 = 7$;
- from L-E2, the optimal action is R, which gives a combined reward of $2 + 1 + 2 = 5$; and
- from R-E4, the optimal action is L, which gives a combined reward of $4 + 3 + 1 = 8$.

Now, from L-E1, we see that R is more profitable, since L-E1-L can at most yield 5, while L-E1-R yields 7. From the root, the value (expected total reward) of L is $\frac{1}{4}(7) + \frac{3}{4}(5) = 5.5$, and the value of R is $\frac{1}{2}(1) + \frac{1}{2}(8) = 4.5$. Thus, the optimal action from the root is L; the expected sum of rewards is 5.5.