Tutorial questions: embedded agents[1]

1. A purely reactive agent is defined by the function $Ag^r : E \to Ac$.

   A standard agent is defined by the function $Ag^s : \mathcal{R}^E \to Ac$.

   (Recall: E is the set of all possible environment states; Ac is the set of all available actions; $\mathcal{R}^E$ is the set of all possible runs over E and Ac that end with an environment state.)

   (a) Is the following statement true? "For every purely reactive agent, there is a behaviourally equivalent standard agent."

   If true, explain how you can define a standard agent that is behaviourally equivalent of any given purely reactive agent.

   If not true, give an example of a purely reactive agent and explain why it is not possible to define a behaviourally equivalent standard agent.

   (b) Is the following statement true? "For every standard agent, there is a behaviourally equivalent purely reactive agent."

   If true, explain how you can define a purely reactive agent that is behaviourally equivalent of any given standard agent.

   If not true, give an example of a standard agent and explain why it is not possible to define a behaviourally equivalent purely reactive agent.

   (Recall: Two agents Ag_1 and Ag_2 are behaviourally equivalent if, for any environment Env, $\mathcal{R}(Ag_1, Env) = \mathcal{R}(Ag_2, Env)$.)

[1] Many thanks to Michael Rovatsos for inspiring these questions.

2. Rock-paper-scissors is a 2-player game where each agent simultaneously selects from one of three moves: rock (R), paper (P), or scissors (S). The outcome is determined as follows:

   - R beats S;
   - S beats P;
   - P beats R;
   - all other combinations result in a draw.

   The set of actions available to the agent are $Ac = \{R, P, S\}$.

   The utility of winning is 1, the utility of losing is -1, and the utility of a draw is 0.

   If we consider a single round of the game we can define the possible environment states as

   $$\{e_0, RR, RS, RP, SR, SS, SP, PR, PS, PP\}$$

   where:

   - e_0 is the initial state before game play has begun, and
   - each two-letter combination XY represents the state in which the agent has played X and its opponent has played Y (e.g., SP means the state in which the agent played scissors and its opponent played paper).

   Assume our opponent (the "environment") plays:

   - P with probability 0.4,
   - S with probability 0.5, and
   - R with probability 0.1.

   Thus for any $\alpha \in Ac$:

   - $\tau((e_0, \alpha)) \in \{RR, RS, RP, SR, SS, SP, PR, PS, PP\}$,
   - the probability that $\tau((e_0, \alpha)) \in \{PP, RP, SP\}$ is 0.4,
   - the probability that $\tau((e_0, \alpha)) \in \{PS, RS, SS\}$ is 0.5, and
   - the probability that $\tau((e_0, \alpha)) \in \{PR, RR, SR\}$ is 0.1.

   We can define a non-deterministic *stochastic* agent

   $$Ag^{st} : \mathcal{R}^E \to 2^{Ac \times [0,1]}$$

   such that $(\alpha, p) \in Ag^{st}(r)$ means that p is the probability the agent will perform action α as the next action of the run r, and the following conditions hold:

   for all $r \in \mathcal{R}^E$:

   - if $(\alpha_i, p_i) \in Ag^{st}(r)$ and $(\alpha_i, p_j) \in Ag^{st}(r)$, then $p_i = p_j$ (i.e., there can be at most one occurrence of each action); and
   - $\sum_{\forall(\alpha_i, p_i) \in Ag^{st}(r)} p_i = 1$ (i.e., the sum of the probabilities associated with all the possible actions must equal 1).

   Specify an *optimal* stochastic agent Ag^{st} for this problem (i.e., one that maximises the expected utility).

Tutorial questions: topic 3

1. Write each of the following statements as a temporal logic formula.

   (a) If a person never dies then they are immortal. (Use the predicates $dies(person)$ and $immortal(person)$.)

   (b) Unless you study agents, you won't know how great they are. (Use the predicates $study(agents, you)$ and $knowGreat(agents, you)$.)

2. The following specification is adapted from a famous "Snow White" example in Concurrent MetateM:

   $$
   \begin{array}{ll}
   snowWhite(ask(X))[give(X)]: & \\
   & \circledcirc ask(X) \rightarrow \Diamond give(X) \\
   & start \rightarrow \Box\neg(give(X) \wedge give(Y) \wedge X \neq Y)
   \end{array}
   $$

   $$
   \begin{array}{ll}
   eager(give(eager))[ask(eager)]: & \\
   & start \rightarrow ask(eager) \\
   & \circledcirc give(eager) \rightarrow ask(eager)
   \end{array}
   $$

   $$
   \begin{array}{ll}
   greedy()[ask(greedy)]: & \\
   & start \rightarrow \Box ask(greedy)
   \end{array}
   $$

   [Note that X and Y are variables. An agent that listens for messages of type $give(X)$ will hear all messages of the form $give(X)$ (where X can be instantiated to anything) (same for an agent who listens for messages of type $ask(X)$). For example, Snow White will listen to messages $ask(eager)$ and $ask(greedy)$.]

   Trace a run of this system (i.e. give a run of the system) for the first 6 time steps (use a table like the one we saw in the example from the core material).

Tutorial questions: topic 4

1. The pseudo-code below defines the control loop for `agent1`, a practical reasoning agent.

```
B := B0;        /* B0 is initial value of B */
I := I0;        /* I0 is initial value of I */
while true do
      get next percept ρ;
      B:= brf(B, ρ);
      D:= options(B, I);
      I:= filter(B, D, I);
      π:= plan(B, I);
      while not (empty(π) or succeeded(I, B)
                  or impossible(I, B)) do
            α:= hd(π);
            execute(α);
            π:= tail(π);
            get next percept ρ;
            B:= brf(B, ρ);
            if reconsider(I, B) then
                 D:= options(B, I);
                 I:= filter(B, D, I);
            end-if
            if not sound(π, I, B) then
                 π:= plan(B, I);
            end-if
       end-while
 end-while
```

   (a) With reference to the pseudo-code that defines `agent1` identify and explain the role of the variables and functions that ensure the *deliberation* of the agent and the *means-ends reasoning* of the agent.

   (b) One of the aims of the pseudo-code that defines `agent1` is to ensure that the agent is committed to plans just as long as it is rational to be committed to them. With reference to the pseudo-code, explain how the control loop achieves this. In your answer, be sure to clearly identify all the circumstances under which an agent reconsiders its plan.

   (c) Another aim of the pseudo-code that defines `agent1` is to ensure that the agent is committed to intentions as long as it is rational to be committed to them. With reference to the pseudo-code, explain how the control loop achieves this. In your answer, be sure to clearly identify all the circumstances under which an agent reconsiders its intentions.

2. Using examples to illustrate your answer, explain how an agent's behaviour is affected by its intentions and the attitudes you can expect an agent to have towards its intentions. The examples you use can be about human agents or artificial agents.

Tutorial questions: topic 5

1. Think of a problem that you think would be well suited to collaborative distributed problem solving.
   (a) Explain how you would decompose the problem into sub-tasks.
   (b) Take one of the sub-tasks from the problem you have described for question 1a and assume that an agent is using the Contract Net protocol to identify a contractor for the sub-task. When making the task announcement, explain
      i. what information the agent would ask to be included in bids and any constraints it would provide in its task announcement; and
      ii. how it would evaluate the bids it receives.
2. Give an example of a pair of actions that should be coordinated (this should be different from the examples in the core material). Explain why these actions should be coordinated.
3. Give an example of prescriptive norm (different from the ones in the core material) and give three different reasons why an agent might choose to violate the norm.

Tutorial questions: topic 6

Consider the following task oriented domain. There are two postal workers $Ag1$ and $Ag2$ who each start at the post office and must deliver some letters to particular houses. So a task relates to a particular letter and is represented by the house it must be delivered to. The cost of a set of tasks is the minimum distance the postal workers must travel to visit each of the houses he/she must deliver to and to return to the post office.

1. A specific scenario is shown in the figure below. We see that there are eight possible houses that the agents may need to visit (the nodes in the figure: $a, b, \ldots, h$) and the post office is at the same location as house a. We assume that the distance between each node is 1 (e.g., the minimum distance from a to b is 1, the minimum distance from a to g is 2).

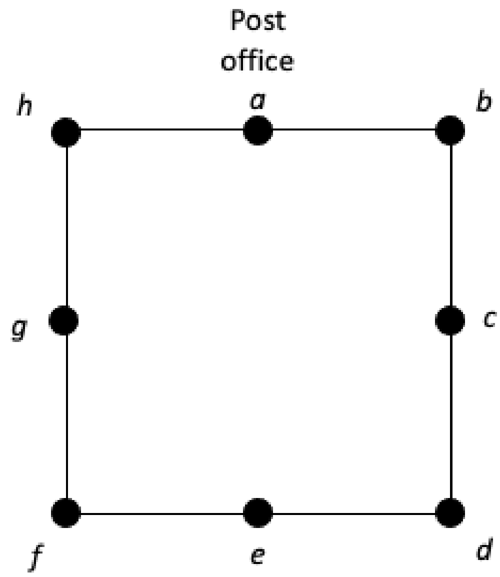

   (a) Consider we have the following initial encounter: $\langle \{b, f\}, \{e\} \rangle$
   So there are three tasks that need performing (to deliver a letter to b, e, and f). Initially, Ag_1 has been allocated the tasks to deliver to both b and f, and Ag_2 has been allocated the task to deliver to e. The cost of this encounter to Ag_1 is 8, and the cost of this encounter to Ag_2 is also 8. (Remember, the cost of a set of tasks is the minimum distance an agent must cover in order to travel from the post office, visit each of the houses, and return to the post office.)

      i. What are all the possible deals for this scenario?

      ii. Which deals are in the negotiation set?

      iii. What happens if the agents each use the Zeuthen strategy with the monotonic concession protocol to try to redistribute their tasks? What is the *expected utility* that each agent can expect to gain from this negotiation? (Reminder: expected utility is equal to the summation over all possible outcomes of the probability that each outcome will occur times the utility of that outcome.) (Reminder: if each agent is equally willing to risk the conflict deal, flip a fair coin to decide who should make a concession.)

(b) Now consider that $Ag1$ deceives $Ag2$ by hiding the fact that it has been allocated the task to deliver to b in the initial allocation.

The (untruthful) initial encounter is thus presented as $\langle\{f\},\{e\}\rangle$, and so the agents will negotiate over the tasks e and f. Whatever the result of the negotiation is, $Ag1$ will also have to perform task b (since this task has been allocated to $Ag1$ in the truthful initial encounter, and $Ag1$ has no chance of giving it away since it has hidden this task from the other agent).

i. Which deals are in the negotiation set, given $Ag1$'s deception here?

ii. What happens if the agents each use the Zeuthen strategy with the monotonic concession protocol to try to redistribute their tasks? What is the expected utility that each agent can expect to gain from this negotiation? (Hint: don't forget that $Ag1$ will also have to perform the task that it hid, to deliver to b.)

2. A different postal office domain scenario is shown in the figure below. Here there are six possible houses that the agents may need to visit. Again, we assume that the distance between each node is 1.

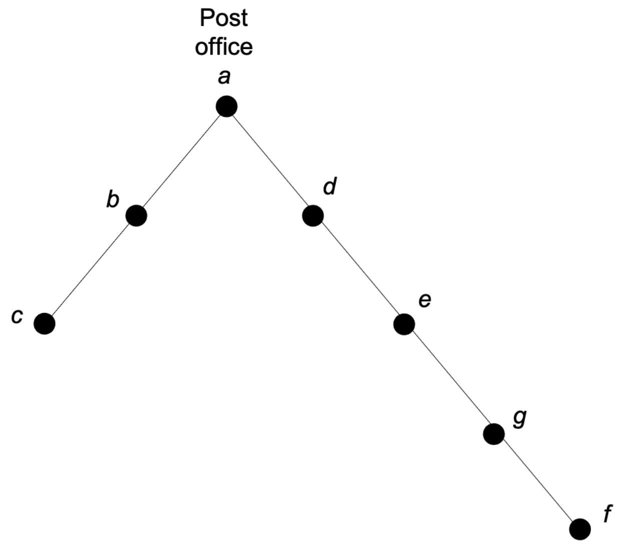

(a) Consider we have the following initial encounter: $\langle\{c, d\}, \{c, d\}\rangle$.

So there are four letters that need to be delivered (two to house c and two to house d) and in the initial allocation each agent has to visit both house c and house d. The cost of this encounter to Ag_1 is 6, and the cost of this encounter to Ag_2 is also 6.

i. Which deals are in the negotiation set?

ii. What happens if the agents each use the Zeuthen strategy with the monotonic concession protocol to try to redistribute their tasks? What is the *expected utility* that each agent can expect to gain from this negotiation? (Reminder: expected utility is equal to the summation over all possible outcomes of the probability that each outcome will occur times the utility of that outcome.) (Reminder: if each agent is equally willing to risk the conflict deal, flip a fair coin to decide who should make a concession.)

(b) Now consider that $Ag1$ deceives $Ag2$ by pretending that it has also been allocated a task to deliver a letter to house f in the initial allocation.

The (untruthful) initial encounter is thus presented as $\langle\{c, d, f\}, \{c, d\}\rangle$.

i. Which deals are in the negotiation set, given $Ag1$'s deception here?

ii. What happens if the agents each use the Zeuthen strategy with the monotonic concession protocol to try to redistribute their tasks? What is the expected utility that each agent can expect to gain from this negotiation? (Hint: whatever the result of the negotiation is, no agent will actually have to visit house f since this task does not actually exist.)

Tutorial questions: making group decisions

1. Consider the following voter profile representing the *truthful* preferences of Alice, Bob, and Charlie:

Alice:	$B \succ A \succ C \succ D$
Bob:	$C \succ A \succ B \succ D$
Charlie:	$B \succ A \succ C \succ D$

   Alice, Bob, and Charlie decide to hold a vote which will be decided according to the *Borda count* protocol, with ties decided randomly.

   Is it in anyone's interest to vote strategically and *falsely* report their preference (assuming they know the truthful preferences of the other two voters)? If your answer is yes, then specify the person as well as the ordering they should report instead of their true preference.

   Hint. A voter v has an incentive to vote strategically if and only if it is the case that:

   - the winner if they submit their true preferences is ω_x
   - the winner if they submit some false preferences is ω_y, and
   - v prefers ω_y to ω_x

2. The plurality voting protocol is an example of a positional scoring rule voting protocol. What is the scoring vector that captures the plurality procedure?

3. Consider the following voter profile:

3 voters:	$A \succ B \succ C$
6 voters:	$B \succ C \succ A$
4 voters:	$C \succ A \succ B$
3 voters:	$C \succ B \succ A$

   Suppose that (s_1, s_2, s_3) is an arbitrary *scoring vector*, where $s_1 \geq s_2 \geq s_3 \geq 0$ and $s_1 > s_3$. Show that the corresponding *positional scoring rule* does not select the Condorcet winner.

   Hint. This is essentially an algebraic manipulation question. You can represent the scores received by each of the three candidates in terms of s_1, s_2 and s_3. Let $score(\omega_i)$ be the score received by candidate ω_i. If you can show that $score(\omega_i) - score(\omega_j) < 0$ then you know that $score(\omega_i) < score(\omega_j)$. If you can show that $score(\omega_i) - score(\omega_j) \leq 0$ then you know that $score(\omega_i) \leq score(\omega_j)$. Don't forget that you know $s_1 \geq s_2 \geq s_3 \geq 0$ and $s_1 > s_3$.

Tutorial questions: allocating scarce resources

1. Auctions on eBay are ascending price open cry auctions but with a proxy bidding system, which proceeds as follows.

   - The auctioneer sets a reserve price (so at the start of the auction the current highest bid is equal to the reserve price), an increment value, and an ending time.
   - Potential buyers use a proxy bidding system. They submit a maximum value to the proxy bidding system (this must be higher than the current highest bid price).
   - Whenever:
     - the current highest bid plus increment is less than or equal to the submitted maximum value, and
     - the buyer is not the current highest bidder,

     the proxy system will bid one increment more than the current highest bid on behalf of the buyer.
   - The potential buyers are able to increase their maximum value at any point during the auction, as long as it is higher than the current highest bid.
   - At the end of the auction, the proxy bidding system ensures that the bidder who submits the highest maximum value wins the auction and the price they pay is the the second highest maximum value submitted plus the increment. If more than one bidder submits the highest maximum value, the bidder who submitted first will win, and the amount they will pay is their maximum value.

   Consider the following three possible bidder strategies.

   Maxing: At the start of the auction, submit to the proxy bidding system a maximum value that is equal to your true valuation of the item and then wait for the auction to close.

   Increments: If you are not the highest bidder and the current highest bid plus increment value is less than your true valuation of the item, submit to the proxy bidding system a maximum value that is the current auction price plus the increment value.

   Sniping: Wait until the last moment before the end of the auction at which it is possible to submit a new maximum value to the proxy bidding system and then submit to the proxy bidding system a maximum value that is equal to your true valuation of the item.

   Assume there are only two potential buyers for an eBay auction, BuyerA and BuyerB. BuyerA's true valuation of the item is £50 and BuyerB's true valuation of the item is £60. The reserve price is £40 and the increment value is £1. What would be the outcome of the auction in the following situations? Say who would win the auction and what price they would end up paying.

   a) BuyerA uses maxing and BuyerB uses maxing.

   b) BuyerA uses increments and BuyerB uses sniping.

   c) BuyerA uses increments and BuyerB uses maxing. (Assume that BuyerA is always first to submit their maximum value.)

   d) BuyerA uses increments and BuyerB uses increments. (Assume that BuyerA is always the first to submit their maximum value.)

   e) BuyerA uses sniping and BuyerB uses maxing.

   f) BuyerA uses sniping and BuyerB uses sniping.

2. Suppose we have two goods a and b and two agents ag_1 and ag_2. The agents declare the following true valuation functions:

$$\hat{v}_1(\{a\}) = v_1(\{a\}) = 8$$

$$\hat{v}_1(\{b\}) = v_1(\{b\}) = 7$$

$$\hat{v}_1(\{a, b\}) = v_1(\{a, b\}) = 14$$

$$\hat{v}_2(\{a\}) = v_2(\{a\}) = 6$$

$$\hat{v}_2(\{b\}) = v_2\{b\}) = 4$$

$$\hat{v}_2(\{a, b\}) = v_2\{a, b\}) = 12$$

a) Which allocation is made using the VCG mechanism?

b) How much does each agent pay to the mechanism?

c) What utility does each agent get?

d) Can ag_2 improve its utility by lying and declaring valuation $\hat{v}_2' \neq v_2$ (assuming ag_1 remains truthful)?

Hint. Think about the possible cases that can occur. In the allocation made when the agents submitted their true valuations we had that ag_2 ends up with nothing, what other possible allocations might have been made and how would these have affected the utility gained by ag_2?

Tutorial questions: reasoning with arguments

1. Below is some knowledge that an agent has about monsters, where X is a variable that can be instantiated by any constant and constants all start with lower-case letters.

   $Angry(cedric)$

   $Strong(cedric)$

   $Sleeping(cedric)$

   $Smiling(cedric)$

   $Smiling(X) \rightarrow \neg Angry(X)$

   $Angry(X) \wedge Strong(X) \rightarrow RunAwayFrom(X)$

   $Sleeping(X) \rightarrow \neg RunAwayFrom(X)$

   (a) Using the knowledge above, it is possible to construct two arguments that **rebut** one another. Identify two such arguments, giving the support and claim of each.

   (b) Using the knowledge above, it is possible to construct two arguments such that one **undercuts** the another. Identify two such arguments, giving the support and the claim of each and making clear which argument undercuts the other.

2. Consider the following abstract argumentation graph.

   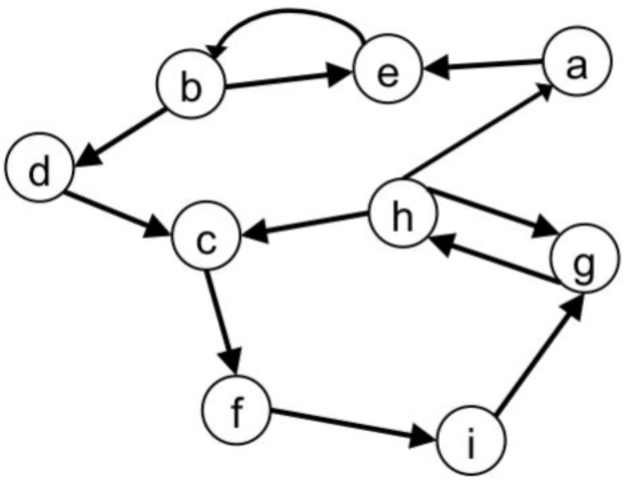

   Which arguments are acceptable under:

   (a) the grounded semantics;

   (b) the preferred credulous semantics;

   (c) the preferred sceptical semantics.

Tutorial questions, topic 10: argument-based dialogues

1. Consider the following simple dialogue system with two participating agents. Each agent has its own knowledge base (its beliefs), which is a finite set of arguments about beliefs.

   Moves:

   An agent can ASSERT an argument.

   An agent can make a CLOSE move.

   Protocol:

   - The agents take it in turn to make a move.
   - The first agent to move ASSERTs an argument, which causes that argument to become the TOPIC of the dialogue.
   - For all other turns:
     - An agent can ASSERT an argument as long as it not previously been asserted and it attacks some argument that has been previously asserted during the dialogue. An agent can always make a CLOSE move.

   Termination rules:

   A dialogue terminates if each agent makes a CLOSE move in succession (i.e. one agent makes a CLOSE move and then the other agent makes a CLOSE move immediately afterwards).

   Dialogue outcome:

   If the TOPIC is acceptable under the grounded semantics given the abstract argumentation framework that is constructed from all of the arguments that have been asserted during the dialogue, then the outcome of the dialogue is Acceptable(TOPIC).

   If the TOPIC is not acceptable under the grounded semantics given the abstract argumentation framework that is constructed from all of the arguments that have been asserted during the dialogue, then the outcome of the dialogue is NotAcceptable(TOPIC).

   Let's consider a few different strategies.

   Strategy 1:

   If it is permissible to ASSERT an argument from your beliefs, then ASSERT some such an argument;

   otherwise make a CLOSE move.

   Strategy 2: If it is permissible to ASSERT an argument from your beliefs and adding that argument to the abstract argumentation framework that is constructed from all off the arguments that have been asserted so far during the dialogue causes the label (i.e. IN, OUT, or UNDEC) of the TOPIC to change, then ASSERT such an argument; otherwise make a CLOSE move.

Strategy 3:

If it is permissible to ASSERT an argument from your beliefs and adding that argument to the abstract argumentation framework that is constructed from all off the arguments that have been asserted so far during the dialogue causes the label of the TOPIC to change to IN, then ASSERT such an argument;

otherwise make a CLOSE move.

Strategy 4:

If it is permissible to ASSERT an argument from your beliefs and adding that argument to the abstract argumentation framework that is constructed from all off the arguments that have been asserted so far during the dialogue causes the label of the TOPIC to change to OUT, then ASSERT such an argument;

otherwise make a CLOSE move.

(a) Assume both of the agents are using Strategy 1. Can you say what type of dialogue (from Walton and Krabbe's dialogue typology) this would produce? If you assume an omniscient perspective where you can see inside the knowledgebase of each agent, can you say what the outcome of the dialogue would be?

(b) Assume both of the agents are using Strategy 2. Can you say what type of dialogue (from Walton and Krabbe's dialogue typology) this would produce? If you assume an omniscient perspective where you can see inside the knowledgebase of each agent, can you say what the outcome of the dialogue would be?

(c) Assume the agent who starts the dialogue was using Strategy 3 and the other agent is using Strategy 4. Can you say what type of dialogue (from Walton and Krabbe's dialogue typology) this would produce?