Foundations of Agents

Example Exam

Markov Decision Processes

question 1 Consider an MDP (S, A, t, r) where $S = \{s, t\}$, $A = \{a, b\}$, and the functions $t : S \times A \times S \to [0, 1]$ and $r : S \times A \times S \to \mathbb{R}$ are specified by the following table:

S	A	S	t	r
s	a	s	0.1	0
s	a	t	0.9	$1\frac{1}{9}$
t	a	s	0.9	0
t	a	t	0.1	0
s	b	s	0.9	2
s	b	t	0.1	2
t	b	s	0.1	3
t	b	t	0.9	3

Determine the optimal policy for this problem using policy iteration starting with the initial policy: $\pi(s) = a$ and $\pi(t) = b$.

Doxastic logic

question 2 Consider the following Kripke structure for two agents: $\mathcal{I} = (S, \pi, R_1^B, R_2^B)$.

Draw an example of a Kripke structure $\mathcal{I}$ for which there holds:
$(\mathcal{I}, s) \not\models B_1 p$
$(\mathcal{I}, s) \models B_2 p$
$(\mathcal{I}, s) \models B_1 \neg B_2 p$

Dynamic logic

question 3 Prove that $\models [a]\varphi \leftrightarrow \neg\langle a\rangle\neg\varphi$ holds.

Dynamic + Epistemic logic

question 4 Consider a simplified wise man puzzle with two wise men, Alex and Bob. The king wishes to know who of his advisors is the wisest. He therefore request the wise men to look at each other. The King tells the wise men that he has three hats, two blue hats and one red hat. Then he tells them that he will place a hat on each of their heads without showing them the color of the hat on their head. Note that Bob can see the hat of Alex and vice versa. The wise men cannot see their own hats.

1. Draw a Kripke structure representing the knowledge of each wise man after the king has placed a hat on each of their heads. The knowledge of each wise man has the properties of S5.
2. Now the king asks Bob whether he knows the color of his hat. Bob answers that he does not.

   Extend the Kripke structure in order to represent this announcement (*the action*) of Bob that he does not know the color of his hat.
3. Next the king asks Alex whether he knows the color of his hat.

   What is Alex's answer. Motivate your answer using the Kripke structure!

BDI logic

question 5 Consider a Kripke structure for describing all information in Bratman's BDI model. The Kripke structure satisfy all relevant axioms. This Kripke structure should represent in a state s that:

- the agent desires $\mathsf{do}(a)$ to hold and desires $\mathsf{do}(b)$ to hold:

$$(\mathcal{I}, s) \models \hat{D}\mathsf{do}(a) \wedge \hat{D}\mathsf{do}(b)$$

- the agent believes that $\mathsf{do}(a) \wedge \mathsf{do}(b)$ can never hold:

$$(\mathcal{I}, s) \models B\Box\neg(\mathsf{do}(a) \wedge \mathsf{do}(b))$$

- the agent intends $\mathsf{do}(b)$ to hold:

$$(\mathcal{I}, s) \models I\mathsf{do}(b)$$

- the agent beliefs that $\mathsf{do}(c; d) \rightarrow \mathsf{do}(b)$ always holds:

$$(\mathcal{I}, s) \models B\Box(\mathsf{do}(c; d) \rightarrow \mathsf{do}(b))$$

Suggestion: do not use more states than necessary.