

Artificial Intelligence 2

Past exam questions

Dr Sean B Holden, January 2013

1 A short historical note...

There is usually some degree of confusion as to precisely which past exam questions are relevant to this course. Allow me to explain. Prior to 2002 this course was essentially what is now AI 1, and at the time AI 1 was a Prolog programming course. The syllabuses changed when I took over the courses in 2002. Consequently, many of the exam questions prior to 2002 are not relevant, although *some* are. The past exam questions mentioned in what follows are the ones that remain relevant to AI 2. Some of the older ones *may* mention subjects not at present in the syllabus. Those items can safely be ignored. Finally, the L^AT_EX files for some of the earlier questions are no longer available so they are not included here. However, copies of the questions can be found in the usual place:

www.cl.cam.ac.uk/teaching/exams/pastpapers/

2 Planning

2012, Paper 8, question 2:

Evil Robot hates kittens. He has invented the *kitty-destroyer* (KD) to rid the world of their menace. To test it, he has a kitten (K) next to the KD on his laboratory bench (B). He has to open the KD, place the kitten inside it, close it and press the start button (SB). He has not however established that this sequence of events will lead to his goal of a destroyed kitten. Evil Robot is equipped with a planning system based on a solver for *constraint satisfaction problems*, and wants to use this to construct a plan.

1. Explain how this problem can be represented using the *state-variable* representation, including in your answer specific examples of a *domain*, a *rigid relation*, a *state variable* and an *action* for the problem. [7 marks]
2. Give *one* reason that a state-variable representation might be preferable to a representation aimed at encoding to a satisfiability problem. [1 marks]
3. Explain, giving a specific example for this problem, how the *action* taken at some time t can be encoded as part of a constraint satisfaction problem. [3 marks]
4. Explain, giving a specific example for this problem, how a *state-variable* can be encoded as part of a constraint satisfaction problem. [4 marks]
5. Explain, giving a specific example for this problem, how a *precondition* for an action can be encoded as part of a constraint satisfaction problem. [5 marks]

2011, Paper 7, question 2:

Consider the following propositional planning problem.

Start state: $\neg A, \neg B, \neg C, D$.

Goal: $A, B, C, \neg D$.

Actions:

- Action 1 has preconditions A, B, C and effect $\neg D$.
- Action 2 has preconditions $\neg A, \neg B$ and effects A and B .
- Action 3 has preconditions $\neg B, \neg C$ and effects B and C .
- Action 4 has precondition B and effect $\neg B$.

1. Using an entire sheet of paper, draw the planning graph as far as state level S_3 , where the start state is at state level S_0 and the first action level is A_0 . Do not add any mutex links at this point. [5 marks]
2. Describe each of the five kinds of *mutex* link that can be incorporated in a planning graph. Add one example of each to the graph you produced in part 1. Clearly label the links to make clear which type they are. [10 marks]
3. At which level in the planning graph will all goals first be present simultaneously? Will the GraphPlan algorithm be able to extract a working plan without extending it beyond this level? Explain your answer, adding further mutex links to your diagram if necessary. [5 marks]

2009, Paper 7, question 4:

Evil Robot has almost completed his Evil Plan for the total destruction of the human race. He has two nasty chemicals, which he has imaginatively called A and B and which are currently stored in containers 1 and 2 respectively. All he has to do now is mix them together in container 3. His designer, an equally evil computer scientist, has equipped Evil Robot with a propositional planning system that allows him to reason about the locations of particular things and about moving a thing from one place to another.


1. Explain how this problem might be represented within a propositional planning system. Give specific examples of the way in which the start state and goal can be represented. [5 marks]
2. Describe in detail an algorithm that can be used to find a plan using this form of representation. [5 marks]
3. Give a specific example of a successor-state axiom using the representation you suggested in part (1). [2 marks]

4. Explain why in this particular planning problem it might be necessary to include one or more precondition axioms and give an example of such an axiom using your representation. [2 marks]
5. Explain why in this particular planning problem it might be necessary to include one or more action exclusion axioms and give an example of such an axiom using your representation. Suggest why it might be unwise to include too many axioms of this type, and explain how a reasonable collection of such axioms might be chosen in a systematic way. [4 marks]
6. Explain how in this problem it might be possible to include state constraints as an alternative to action exclusion axioms, and give a specific example of such a constraint using your representation. [2 marks]

2008, Paper 7, question 6:

We have a simple, propositionalised planning problem and we suspect that we might be able to solve it using the *GraphPlan* algorithm. The problem is as follows.

Action *A* has preconditions $\{\neg X\}$ and effects $\{X, Z\}$, action *B* has preconditions $\{\neg Y\}$ and effects $\{X, Y\}$, and action *C* has preconditions $\{X, Y, Z\}$ and effects $\{W\}$. The start state for the problem is $\{\neg W, \neg X, \neg Y, \neg Z\}$ and the goal is $\{W\}$.




1. Labelling the start state as level S_0 and the first action level as A_0 , draw the planning graph for this problem up to and including level S_2 . Use an entire sheet of paper for this diagram. [5 marks]
2. Describe each of the five kinds of *mutex link* that can appear in a planning graph, and add an example of each to the diagram drawn in part (a), clearly labelling it to show which kind of mutex link it is. [10 marks]
3. What is the *level cost* of a literal in a planning graph? Explain why this measure of cost might perform poorly as a measure of how hard the literal is to achieve, and suggest a way in which its performance might be improved. [2 marks]
4. Will GraphPlan be able to extract a working plan from the diagram you have drawn in parts (a) and (b)? Explain your answer. You may if you wish add further mutex links to your diagram at this stage. [3 marks]

3 Bayesian Networks

2009, Paper 8, question 1:

Consider the following Bayesian Network:




The associated probability distributions for the binary random variables A , B , C and D are $\Pr(a) = 0.7$, $\Pr(\neg a) = 0.3$ and:

A	$\Pr(b A)$	B	$\Pr(c B)$	B	C	$\Pr(d B, C)$
\top	0.1	\top	0.2	\top	\top	0.6
\perp	0.15	\perp	0.95	\top	\perp	0.5
				\perp	\top	0.4
				\perp	\perp	0.3

1. Write down an expression for the full joint distribution of the random variables A , B , C and D . Compute the probability that A and B are \top while C and D are \perp . [2 marks]
2. Use the variable elimination algorithm to compute the probability distribution of B conditional on the evidence that $D = \perp$. [16 marks]
3. Explain why the variable elimination might not be an effective algorithm to use in practice and suggest an alternative that addresses the shortcoming you have given. [2 marks]

2006, paper 8, question 9:

Consider the following Bayesian network:



The associated probability distributions for the binary random variables A , B , C and D are $\Pr(a) = 0.1$, $\Pr(\neg a) = 0.9$, $\Pr(b) = 0.8$, $\Pr(\neg b) = 0.2$, and:

A	B	$\Pr(c A, B)$	B	C	$\Pr(d B, C)$
\top	\top	0.5	\top	\top	0.2
\top	\perp	0.6	\top	\perp	0.9
\perp	\top	0.8	\perp	\top	0.8
\perp	\perp	0.7	\perp	\perp	0.1

1. Explain why the representation of the joint distribution of A , B , C and D using the Bayesian network is preferable to a direct tabular representation. [2 marks]
2. Use the *variable elimination algorithm* to compute the probability distribution of B conditional on the evidence that $D = \top$. [16 marks]
3. Comment on the computational complexity of the variable elimination algorithm. [2 marks]

2005, paper 8, question 2:

1. A given probabilistic inference problem involves a query random variable (RV) Q , evidence RVs $\mathbf{E} = (E_1, \dots, E_n)$ and unobserved RVs $\mathbf{U} = (U_1, \dots, U_m)$. Assuming that RVs are discrete, state the equation allowing the inference $\Pr(Q|\mathbf{E} = (e_1, \dots, e_n))$ to be computed using the full joint distribution of the RVs and explain why in practice such a method might fail. [5 marks]
2. Give a general definition of a *Bayesian network* (BN), and explain how a BN represents a joint probability distribution. [4 marks]
3. Define *conditional independence* and explain how BNs make use of this concept to reduce the effect of the difficulties mentioned in your answer to (a). Describe the way in which conditional independence is employed by the *naive Bayes* algorithm. [6 marks]
4. Describe two further issues relevant to the application of BNs in a practical context and describe briefly how these issues can be addressed. [5 marks]

4 Value of Perfect Information

2011, paper 8, question 8:

1. Give a definition of *expected utility* and explain why the concept is useful in the context of decision-making. [2 marks]
2. Give a definition of the *value of perfect information* and explain why the concept is useful in the context of decision-making. [4 marks]

3. A talented, but nervous, student has to sit a difficult and important examination. There are only two possible outcomes: `pass` or `fail` and the student attaches to these utilities of $U(\text{pass}) = 10^6$ and $U(\text{fail}) = 10^8$. Lacking in confidence, his beliefs are that $\Pr(\text{pass}|\text{revise}) = 0.55$ and $\Pr(\text{pass}|\neg\text{revise}) = 0.2$. Calculate the expected utility of the situation described. [4 marks]
4. The student finds what he believes might be a copy of this years examination paper, discarded by a careless examiner. He believes that

$$\Pr(\text{pass}|\text{revise}, \text{thisYearsPaper}) = 0.75$$

However, should he be wrong then

$$\Pr(\text{pass}|\text{revise}, \neg\text{thisYearsPaper}) = 0.1$$

as he will waste time learning to answer the wrong questions, because he will revise from the wrong paper. Not revising implies

$$\Pr(\text{pass}|\neg\text{revise}, \text{thisYearsPaper}) = 0.7$$

However, should he be wrong then

$$\Pr(\text{pass}|\neg\text{revise}, \neg\text{thisYearsPaper}) = 0.08$$

He considers bribing somebody to tell him whether he has this years paper or not; however, he thinks it is unlikely that he in fact has this years paper, and therefore believes that

$$\Pr(\text{thisYearsPaper}) = 0.7$$

Compute the value of perfect information associated with finding out whether the paper is the right one. [10 marks]

2007, paper 8, question 9:

An agent can exist in a state $s \in S$ and can move between states by performing actions, the outcome of which might be uncertain.

1. Explain what is meant by a *Utility Function* within this context. [2 marks]
2. Give a definition of *Maximum Expected Utility* and describe the way in which it can be used to decide which action to perform next. [3 marks]
3. What difficulties might you expect to have to overcome in practice in order to implement such a scheme? [3 marks]
4. Explain why it makes sense to use a utility function in the design of an agent, even though it can be argued that real agents (such as humans) appear not to do this, but rather to act on the basis of *preferences*. [4 marks]
5. As well as actions allowing an agent to move between states, an agent might be capable of performing actions that allow it to discover more about its environment. Give a full derivation of the *Value of Perfect Information*, and explain how this idea can be used as the basis for an agent that can gather further information in a way that takes account of the potential cost of performing such actions. [8 marks]

5 Machine Learning

2010, Paper 8, question 2:

Consider the following learning problem in which we wish to classify inputs, each consisting of a single real number, into one of two possible classes C_1 and C_2 . There are three potential hypotheses where $\Pr(h_1) = 3/10$, $\Pr(h_2) = 5/10$ and $\Pr(h_3) = 2/10$. The hypotheses are the following functions

$$h_i(x) = x - \frac{i-1}{5}$$

and the likelihood for any hypothesis h_i is

$$\Pr(x \in C_1 | h_i, x) = \sigma(h_i(x))$$

where $\sigma(y) = 1/(1 + \exp(-y))$. You have seen three examples: $(0.9, C_1)$, $(0.95, C_2)$ and $(1.3, C_2)$, and you now wish to classify the new point $x = 1.1$.

1. Explain how in general the *maximum a posteriori* (MAP) classifier works. [3 marks]
2. Compute the class that the MAP classifier would predict in this case. [10 marks]
3. The preferred alternative to the MAP classifier is the Bayesian classifier, computing $\Pr(x \in C_1 | x, \mathbf{s})$. where \mathbf{s} is the vector of examples. Show that

$$\Pr(x \in C_1 | x, \mathbf{s}) = \sum_{h_i} \Pr(x \in C_1 | h_i, x) \Pr(h_i | \mathbf{s})$$

What are you assuming about independence in deriving this result? [3 marks]

4. Compute the class that the Bayesian classifier would predict in this case. [4 marks]

6 Hidden Markov Models

2008, paper 9, question 5:

A friend of mine likes to climb on the roofs of Cambridge. To make a good start to the coming week, he climbs on a Sunday with probability 0.98. Being concerned for his own safety, he is less likely to climb today if he climbed yesterday, so

$$\Pr(\text{climb_today} | \text{climb_yesterday}) = 0.4$$

If he did not climb yesterday then he is very likely to climb today, so

$$\Pr(\text{climb_today} | \neg \text{climb_yesterday}) = 0.1$$

Unfortunately, he is not a very good climber, and is quite likely to injure himself if he goes climbing, so

$$\Pr(\text{injury} | \text{climb_today}) = 0.8$$

whereas

$$\Pr(\text{injury} | \neg \text{climb_today}) = 0.1$$

1. Explain how my friend's behaviour can be formulated as a *Hidden Markov Model*. What assumptions are required? [4 marks]
2. You learn that on Monday and Tuesday evening he obtains an injury, but on Wednesday evening he does not. Use the *filtering* algorithm to compute the probability that he climbed on Wednesday. [8 marks]
3. Over the course of the week, you also learn that he does not obtain an injury on Thursday or Friday. Use the *smoothing* algorithm to compute the probability that he climbed on Thursday. [8 marks]

2010, Paper 7, question 4:

Professor Elbow-Patch is not the man he used to be, and in particular has a tendency to fall over for no apparent reason. This problem is made worse if he has drunk port with his dinner. He almost always drinks port on a Sunday, and if he drinks on any given day he is unlikely—for the sake of his long-suffering liver—to drink port on the following day. However, if he does not drink on a given day then he is very likely to succumb to temptation on the following day.

The probability that he falls over after drinking is $\Pr(\text{fall}|\text{drank}) = 0.7$. The probability that he falls over when he hasn't drunk is $\Pr(\text{fall}|\neg\text{drank}) = 0.1$. He drinks on a Sunday with probability 0.9. If he has not drunk on a given day then the probability that he drinks the following day is $\Pr(\text{drink today}|\neg\text{drank yesterday}) = 0.8$. If he has drunk on a given day then the probability that he drinks the following day is $\Pr(\text{drink today}|\text{drank yesterday}) = 0.1$.

1. Explain how this problem can be represented as a *hidden Markov model*. What assumptions are required? [4 marks]
2. Denoting observations at time i by E_i and states at time i by S_i give a derivation of the *filtering algorithm* for computing $\Pr(S_t|E_1, \dots, E_t)$. [8 marks]
3. You observe the Professor on Sunday, Monday and Tuesday and notice that he doesn't fall over at all. Use the filtering algorithm to compute the probability that he drank port on Tuesday. [8 marks]

2005, paper 9, question 8:

We wish to model the unobservable state of an environment using a sequence $S_0 \rightarrow S_1 \rightarrow S_2 \rightarrow \dots$ of sets of random variables (RVs) where at time i we are in state S_i and observe a set of RVs E_i . The distributions of the RVs do not change over time, and observations depend only on the current state.

1. Define a *Markov process*, the *transition model* and the *sensor model* within this context. [3 marks]
2. Assuming that evidence $E_{1:t} = e_{1:t} = (e_1, e_2, \dots, e_t)$ has been observed define the tasks of *filtering*, *prediction* and *smoothing*. [3 marks]
3. Derive a recursive estimation algorithm for performing filtering by combining the evidence e_t obtained at time t with the result of filtering at time $t - 1$. [8 marks]

4. How does a *hidden Markov model* differ from the setup described? [1 mark]
5. Show how for the case of a hidden Markov model your filtering algorithm can be expressed in terms only of matrix operations. [5 marks]


7 Reinforcement Learning

2012, Paper 7, question 2:

A *reinforcement learning problem* has states $\{s_1, \dots, s_n\}$, actions $\{a_1, \dots, a_m\}$, reward function $R(s, a)$ and next state function $S(s, a)$.

1. Give a general definition of a *policy* for such a problem. [1 mark]
2. Give a general definition of the *discounted cumulative reward* and the corresponding *optimal policy* for such a problem. [5 marks]
3. Give an expression for the optimal policy in terms of R , S and the discounted cumulative reward, and show how this can be modified to produce the *Q-learning algorithm*. [7 marks]

In a simple reinforcement learning problem, states are positions on a grid and actions are up and right. The only way an agent can receive a reward is by moving into one of two special positions, one of which has a reward of 10 and the other of -100 .



Here, states are labelled by their grid coordinates. A possible sequence of actions (sequence 1) is shown by solid arrows, ending with a reward of 10 being received, and another (sequence 2) by dashed arrows ending with a reward of -100 .

4. Assume that all Q values are initialised at 0.
 - (a) Explain how the Q values are altered if sequence 1 is used *twice* in succession by the Q -learning algorithm. [4 marks]

- (b) Explain what further changes occur to the Q values if sequence 2 is then used *once* by the Q -learning algorithm. [3 marks]

2007, paper 9, question 9:

An agent exists within an environment in which it can perform actions to move between states. On executing any action it moves to a new state and receives a reward. The agent aims to explore its environment in such a way as to learn which action to perform in any given state so as in some sense to maximise the accumulated reward it receives over time.

1. Give a detailed definition of a *deterministic Markov decision process* within the stated framework. [4 marks]
2. Give a general definition of a *policy*, of the *discounted cumulative reward*, and of the *optimum policy* within this framework. [4 marks]
3. Give a detailed derivation of the Q -learning algorithm for learning the optimum policy. [8 marks]
4. Explain why it is necessary to trade-off *exploration* against *exploitation* when applying Q -learning, and explain one way in which this can be achieved in practice. [4 marks]