

Symbolic Machine Learning - Model Exam Questions

PUBLISHED SPECIMEN

Question 1. (10 points)

We have a bag of three biased coins a , b , and c with probabilities of coming up heads of 20%, 60%, and 80%, respectively. One coin is drawn randomly from the bag (a is twice more likely to be drawn than the other two coins), and then the coin is flipped three times to generate the outcomes X_1 , X_2 , and X_3 .

1. (4 points) Draw the Bayesian network corresponding to this setup and define the necessary CPTs.
2. (6 points) Calculate which coin was most likely to have been drawn from the bag if the observed flips come out heads twice and tails once.

PUBLISHED SPECIMEN

Question 2. (15 points)

Design an on-line learning agent that learns an s -DNF (i.e., a disjunction of conjunctive terms, where each term has at most s literals) on n propositional variables from only *negative observations* (i.e, truth-assignments for which the DNF is not satisfied).

1. (5 points) Describe the initial hypothesized DNF (what terms will be in it?) and describe how the DNF is updated with each negative example.
2. (5 points) Provide a bound, as tight as you can, on the maximum number of errors the agents makes, without any assumption on the set of observations received. Is the bound polynomial or exponential in (i) s , (ii) n ?
3. (5 points) Assume $n = s = 2$. Demonstrate the sequence of hypothesized DNF's, starting with the initial one and updating on two successive observations $o_1 = (1, 0)$, $o_2 = (0, 1)$.

PUBLISHED SPECIMEN

Question 3. (5 points)

Consider a version-space agent whose version space contains all non-contradictory conjunctions on 3 propositional variables.

1. (3 points) How large will the initial (i.e., before seeing any observation) version space be? Explain your reasoning.
2. (2 points) What is the maximal possible size of it after the first update, assuming the first class decision was wrong? Explain your reasoning.

PUBLISHED SPECIMEN

Question 4. (10 points)

Determine the least general generalization of the following two assertions


1. *Superman is mortal or he is not a human.*
2. *Every human who smokes is mortal.*

by representing them as first-order logic clauses and computing their least general generalization with respect to the θ -subsumption order.


PUBLISHED SPECIMEN

Question 5. (10 points)

Consider the deterministic world below (part (a)). Allowable moves are shown by arrows, and the numbers indicate the reward for performing each action. If there is no number, the reward is zero. Given the Q values in (b), show the changes in the Q estimates when the agent takes the path shown by the dotted line (the agent starts in the lower left cell) when $\gamma = 0.5$ and $\alpha = 0.5$. Show all of your work including intermediate steps.



(a)



(b)