CS781 End-semester Exam (Autumn 2024)

Max marks: 55 Duration: 180 mins

- *The exam is open book and notes. However, you are not allowed to search on the internet or consult others over the internet for your answers.*
- *Be brief, complete and stick to what has been asked.*
- *Unless asked for explicitly, you may cite results/proofs covered in class without reproducing them.*
- ***If you need to make any assumptions, state them clearly.***
- ***Do not copy solutions from others. Penalty for offenders: FR grade.***

1. *[15 marks]*

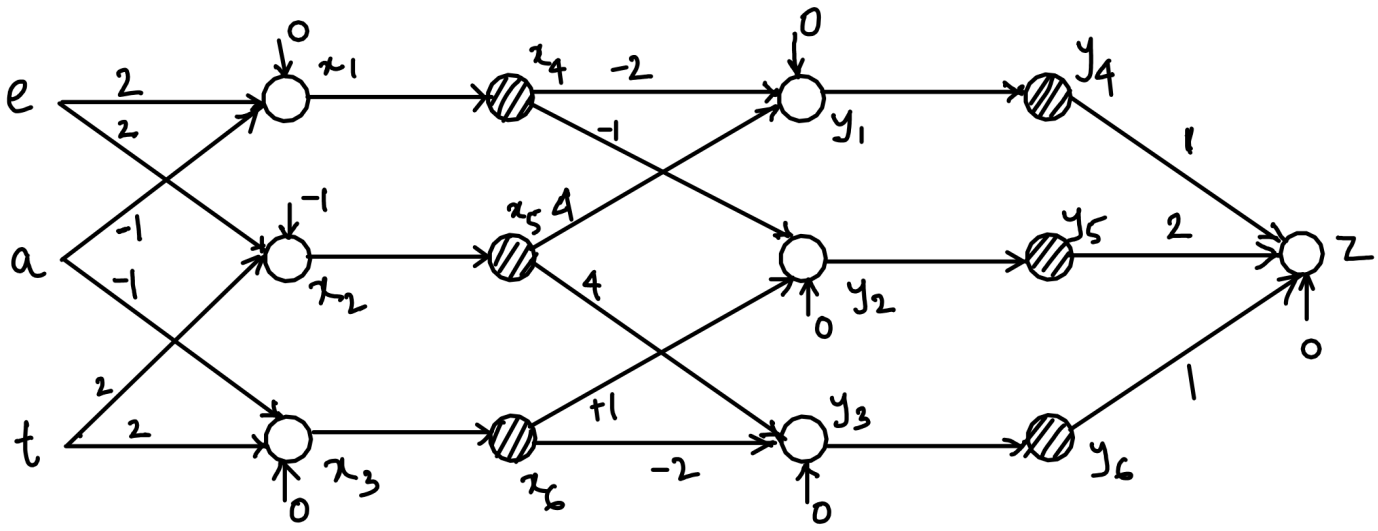

Figure 1: Neural network for admission application processing

The neural network shown in Fig. 1 is used to determine if an admission application meets the bar for review by an admissions officer. The network has three inputs: e denotes normalized exam score, a denotes normalized age of the applicant, and t denotes normalized transcript score. All linear aggregation nodes (unshaded) have bias as indicated, and all shaded nodes represent ReLUs. The three normalized inputs to the network have range $[0, 1]$. An application with normalized score (e, a, t) is said to *meet the bar* if the output z has value ≥ 9.

(a) *[2 marks]* An application has the following normalized input values $e = 0.8, a = 0.5, t = 0.5$. Does this application *meet the bar* by the criteria given above?

(b) We wish to find if the answer to the above question (i.e. does the application meet the bar) stays unchanged if the e, a, t values of the application are changed slightly. Towards this end, we consider $e \in [0.7, 0.9]$, $a \in [0.3, 0.6]$ and $t \in [0.3, 0.6]$.

i. *[13 × 0.5 = 6.5 marks]* Using ***only interval propagation and the fact that the output of a ReLU is non-negative***, find the missing intervals for the variables in the following table. Please refer to Fig. 1 to understand what each variable denotes.

e	$\in$	[0.7, 0.9]	a	$\in$	[0.3, 0.6]
t	$\in$	[0.3, 0.6]	x_1	$\in$	
x_2	$\in$		x_3	$\in$	
x_4	$\in$		x_5	$\in$	
x_6	$\in$		y_1	$\in$	
y_2	$\in$		y_3	$\in$	
y_4	$\in$		y_5	$\in$	
y_6	$\in$		z	$\in$	

ii. *[3.5 marks]* A student claims that the output z of the neural network in Fig. 1 can be expressed as a linear function of of e, a, t, when $e \in [0.7, 0.9], a \in [0.3, 0.6], t \in [0.3, 0.6]$? In other words, the student claims that $z = P.e + Q.a + R.t + S$, where P, Q, R, S are real-valued constants, when e, a, t are in the given ranges.
If you think the student is correct, give values of P, Q, R, S along with proper justification. Otherwise, argue why no real values of P, Q, R, S suffice for the student's claim.
Answers without justification will fetch no marks.

iii. *[1.5 × 2 marks]* Use your answer to the previous question to give real values $e_1, a_1, t_1, e_2, a_2, t_2$ such that the following three conditions are satisfied: (i) $(e_1, a_1, t_1) \neq (0.8, 0.5, 0.5)$, (ii) $(e_2, a_2, t_2) \neq (0.8, 0.5, 0.50)$, (iii) an application with (e_1, a_1, t_1) meets the bar, and (iv) an application with (e_2, a_2, t_2) does not meet the bar. See the initial part of the question to understand when an application meets the bar.

2. *[10 marks]* Consider the 1-input neural network shown in Fig. 2. The shaded nodes represent ReLUs, while unshaded nodes are simply linear aggregators (possibly with biases). The input x_0 is assumed to take values in the range $[-3, 3]$. The intervals corresponding to some internal nodes are as shown in the figure, and these may be assumed to be correct without justification.

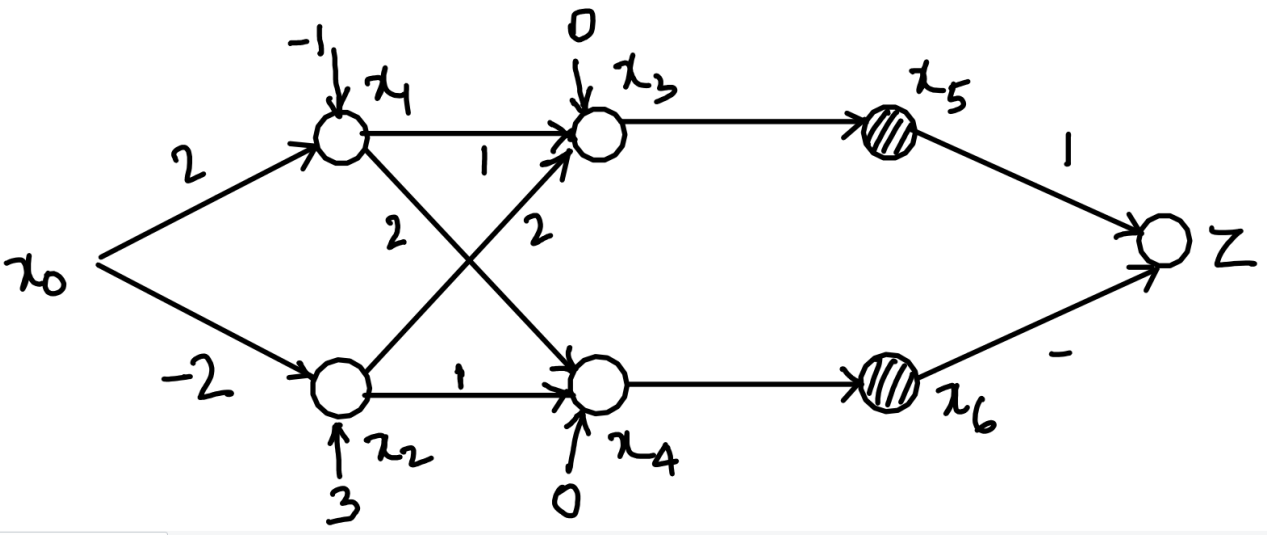

Figure 2: Neural network with single input

(a) *[5 marks]* We wish to use the technique of α-CROWN to find lower and upper bound expressions of z in terms of the input x_0 and some α_i parameters. Specifically, we will use α_1 $(0 \leq \alpha_1 \leq 1)$ to construct linear relaxations of $x_5 = ReLU(x_3)$. Similarly, we will use α_2 $(0 \leq \alpha_2 \leq 1)$ to construct linear relaxations of $x_6 = ReLU(x_4)$. Using these parameters, find the best linear upper and lower bound expressions for z in terms of x_0, where the coefficients in the linear expressions can be functions of α_1, α_2.

You must show all steps in the derivation of your expressions.

(b) *[5 marks]* To find as good an upper bound of z as we can, we must choose values of α_1, α_2 such that $0 \leq \alpha_1, \alpha_2 \leq 1$ and the upper bound expression is minimized as x_0 varies in $[-3, 3]$. Similarly, to find as good a lower bound of z as we can, we must choose α_1, α_2 such that the lower bound expression is maximized as x_0 varies in $[-3, 3]$. Note that the values of α_1, α_2 chosen for maximizing the lower bound expression need not be the values chosen for minimizing the upper bound expression.

Find as good concrete upper and lower bounds of z as you can, and the corresponding values of α_1 and α_2, following the approach outlined above. You must justify your choices of α_1, α_2 in each case. Simply providing the values will fetch no marks.

3. *[10 marks]* Consider the Markov decision process (MDP) and safety automaton shown in Fig. **??**.

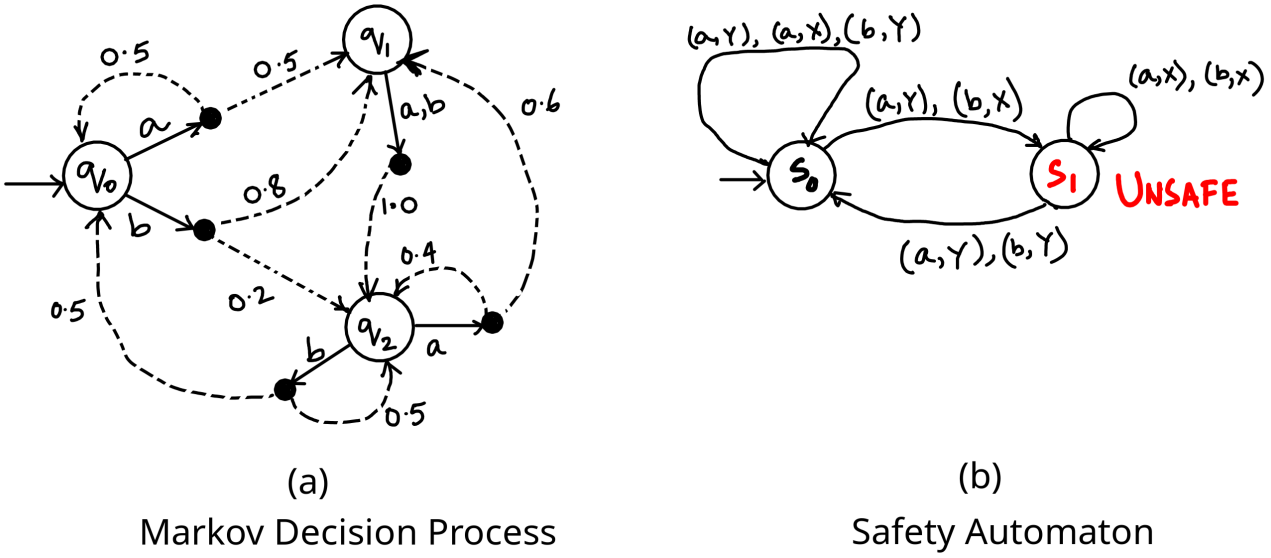

(a) Markov Decision Process

(b) Safety Automaton

Figure 3: MDP and safety automaton

These are the same MDP and safety automaton as you saw in Quiz 2. Assume the observation function for MDP states is given by $f : \{q_0, q_1, q_2\} \to \{X, Y\}$, where $f(q_0) = f(q_2) = X$ and $f(q_1) = Y$. When the MDP is in state q_i and the agent takes action $\alpha \in \{a, b\}$, the safety automaton makes a non-deterministic transition on the input $(\alpha, f(q_i))$. For example, if the sequence of states and actions of the MDP is $q_0\ a\ q_1\ a\ q_2\ a\ q_2\ a \ldots$, the safety automaton moves on the sequence of inputs $(a, f(q_0)), (a, f(q_1), (a, f(q_2)), (a, f(q_2)), \ldots$, i.e. on the input sequence $(a, X), (a, Y), (a, X), (a, X), \ldots$. Note that the automaton is non-deterministic, so it can have multiple runs on the above input sequence. State S_1 is the unsafe state in the automaton. As long as the automaton has even a single run that encounters the unsafe state, we consider the interaction of the agent and environment to be unsafe.

In Quiz 2, you were asked to design a pre-emptive shield for the agent. In this question, we wish to design a *post-posed shield.*

Use the non-deterministic automaton resulting from ignoring the probabilities on the dashed edges of the MDP as the abstraction of the MDP. You need not draw the state corresponding to infeasible runs of the MDP.

(a) *[5 marks]* Construct the product automaton (game graph) and identify all product states that are in the winning region.

(b) *[5 marks]* Construct a post-posed shield for the agent, using the game graph obtained above.

4. *[10 marks]* An input to a black-box ML model has a large number (say, $100,000$) of binary (0/1 valued) features, say $f_1, \ldots f_N$. On being fed such an input, the model responds with a 0/1 answer. An ML engineer suspects that that the classification task being done by the black-box model, can also be reasonably (if not exactly) done by inspecting a single feature, but she is unsure which feature this should be. So, she sets out to use some of the techniques studied in class to symbolically search over the very large space of candidate interpretations.

   Specifically, each hypothesis (interpretation) in our hypothesis class is a decision tree (DT) consisting of one internal node (also its root) labeled by one of the features in $\{f_1, \ldots f_N\}$ and two leaves, one of which is labeled 0 and the other labeled 1. Note that there are exactly two such DTs for every f_r, since the leaf labeled 0 could be reached when f_r is 0, or when f_r is 1. We ignore trees where the same leaf is reached regardless of the value of f_r, as these represent trivial DTs,

   Of course, the engineer needs to sample the input-output space of the black-box a minimum number of times (given by the sample complexity of PAC learning). Let us say this sample count is M and the sample set is S. Each sample s_i $(1 \leq i \leq M)$ in S is therefore a tuple $(f_{1,i}, \ldots f_{N,i}, l_i)$, where $f_{r,i}$ is the value (0/1) of feature f_r in sample s_i, and l_i is the output of the black-box model for sample s_i.

   Let p_i and q_i be 0/1 integer variables, for $1 \leq i \leq N$, with the following meanings:

   - p_i is 1 iff the DT has its root labeled f_i and the 0 (resp. 1) leaf is reached when f_i is 0 (resp. 1).
   - q_i is 1 iff the DT has its root labeled f_i and the 1 (resp. 0) leaf is reached whern f_i is 0 (resp. 1).

   (a) *[5 marks]* Write a (mixed) integer linear program in terms of p_i's and q_i's such that a solution to the program gives the "most accurate" DT in our hypothesis class for the sample set S. For purposes of this question, DT T_1 is considered more accurate than DT T_2 if T_1 correctly predicts the outcome of more samples in S than T_2 does.

   If you need to use additional variables beyond the p_i's and q_i's, you must clearly indicate what each variable stands for, similar to what has been done above for p_i and q_i. You must clearly indicate the *linear* constraints in your (mixed) integer linear program, and also the *linear* objective function that you wish to maximize.

   (b) *[3 marks]* Suppose each f_i has a "desirability" score, given by a real number $d_i \in (0, 1]$. Thus, a hypothesis tree with its root labeled by f_i is more desirable than one with its root labeled f_j iff $\frac{d_i}{d_i+d_j} > \frac{d_j}{d_i+d_j}$. With the desirability score, we can now talk of Pareto-optimal (between accuracy and desirability) DTs. We wish to find (any) one such Pareto-optimal DT by using the same (mixed) integer linear program as in the previous sub-question, but by modifying the linear objective function appropriately.

   Indicate what linear objective function you will use in this case, along with justification.

   (c) *[2 marks]* Let A denote the accuracy score of a DT, and D denote its desirability score. Each of A and D range over $[0, 1]$ in general. After some investigation, the engineer makes the following observations:

   - There are several DTs with $A = 0.5$ that maximize the value of $A + 2D$, when considering all DTs with $A \in [0, 1]$.
   - There are several DTs with $A = 0.2$ that maximize the value of $A + D$ when considering all DTs with $A \in [0, 0.5)$
   - There is one DT with $A = 0.8$ that maximizes the value of $A + 0.5D$ when considering all DTs with $A \in (0.5, 1]$
   - There are no DTs with $A \in [0, 0.2) \cup (0.2, 0.5) \cup (0.5, 0.8) \cup (0.8, 1]$.

Indicate the maximum and minimum number of Pareto-optimal points in the accuracy vs desirability plot for DTs from our hypothesis class. Answers without justification will fetch no marks.

5. *[10 marks]* The interaction of an environment with a reinforcement learning agent that has learnt a deterministic strategy, can often be described as a Markov chain. Consider such a Markov chain with 5 states $\{s_1, s_2, s_3, s_4, s_5\}$, but with an unknown transition matrix M. Suppose the interaction modeled by the Markov Chain is being monitored by a frequentist monitor. Every time the Markov chain reaches a state s_i, the monitor observes an event i.

   Now consider the following sequence of observations of the monitor, where repeats are denoted in parentheses (e.g., 1(5) 2 3(2) denotes the sequence 1 1 1 1 1 2 3 3)

   1 2(3) 1(2) 2(2) 3(2) 4 2 3(3) 2(3) 5 4(3) 2 5(2) 1(3) 5(2) 2 3 4(3) 2 5(3) 1(4) 2 3(3) 2(2) 3

   Let $M_{i,j}$ denote the probability of the Markov chain transitioning to state s_j from state s_i.

   For each of the following transition probability matrix entries or functions of such entries, give the 90% confidence interval from the sequence given above. Thus, when estimating $M_{i,j}$, you must give an interval $[l_{i,j}, u_{i,j}]$ such that $\mathsf{Pr}\big(l_{i,j} \leq M_{i,j} \leq u_{i,j}\big) \geq 0.9$

   (a) *[3 marks]* $M_{4,4}$

   (b) *[3 marks]* $M_{1,1} + M_{2,3}$

   (c) *[4 marks]* $M_{2,3} \times M_{2,5}$.