COMP7015 Artificial Intelligence **Semester 1, 2022/23**

Written Assignment 2

Instructor: Dr. Kejing Yin Nov. 18, 2022 (due: 23:59 pm, Nov. 30)

Problem 1: Computational Graphs and Backpropagation (30 marks)

Given a sequential dataset $\mathcal{D} = \{(\mathbf{x}_1, y_1), \ldots, (\mathbf{x}_T, y_T)\}$, consider a vanilla RNN model, which computes the loss function at the t-th time step using the following formulas:

$$\begin{aligned}\mathbf{a}_t =& \mathbf{b} + \mathbf{W}\mathbf{h}_{t-1} + \mathbf{U}\mathbf{x}_t,\\ \mathbf{h}_t =& \tanh(\mathbf{a}_t),\\ o_t =& c + \mathbf{v}^\top \mathbf{h}_t,\\ L_t =& (o_t - y_t)^2,\end{aligned}$$

where $\tanh(z) = \frac{e^z - e^{-z}}{e^z + e^{-z}}$ and the final loss function is given by $\text{TrainLoss} = \sum_{t=1}^{T} L_t$ with T being the length of the sequence. For simplicity, let's assume the input data is three dimensional (*i.e.*, $\mathbf{x} \in \mathbb{R}^3$), the hidden size (the dimension of $\mathbf{h}$) is 2 and $T = 2$ in this problem.

Q1. Construct the computational graph.

Q2. What are the model parameters to be learned?

Q3. Given the following model parameters and data, compute the loss function in forward pass and compute the gradient with respect to all model parameters in backward pass. We follow the common practice to use the zero initial hidden states, *i.e.*, $\mathbf{h}_0 = [0, 0]^\top$.

$$\begin{aligned}\mathbf{b} =& [0.1, 0]^\top\\ \mathbf{W} =& \begin{bmatrix} 0.5 & -0.4 \\ -0.3 & 0.7 \end{bmatrix}\\ \mathbf{U} =& \begin{bmatrix} 0.8 & -0.3 & 0.5 \\ -1 & 0.1 & 0.9 \end{bmatrix}\\ c =& 0.7\\ \mathbf{v} =& [0.5, 0.8]^\top\\ \mathcal{D} =& \{([1, -1, 1]^\top, 1.5), ([1, -1, -1]^\top, 0.8)\}\end{aligned}$$

Problem 2: Naïve Bayes Classifier (25 marks)

Given the following dataset about watermelons:

Features				Label
color	root	texture	surface	ripe
green	curly	clear	hard	yes
dark	curly	clear	hard	yes
dark	curly	clear	hard	yes
green	curly	clear	hard	yes
light	curly	clear	hard	yes
green	slightly curly	clear	soft	yes
dark	slightly curly	slightly blurry	soft	yes
dark	slightly curly	clear	hard	yes
dark	slightly curly	slightly blurry	hard	no
green	straight	clear	soft	no
light	straight	blurry	hard	no
light	curly	blurry	soft	no
green	slightly curly	slightly blurry	hard	no
light	slightly curly	slightly blurry	hard	no
dark	slightly curly	clear	soft	no
light	curly	blurry	hard	no
green	curly	slightly blurry	hard	no

What would be the prediction generated by a naïve Bayes classifier that is trained using the above dataset for a new example that have features: color=green, root=slightly curly, texture=clear, and surface=hard? Show the steps of how you get your results.

Problem 3: Bayesian Networks (25 marks)

Answer the following questions.

Q1. Draw out the Bayesian Network that corresponds to the following factorization of joint probability.

(a) $P(A)P(B)P(C)P(D)$

(b) $P(D|A,B,C)P(A)P(B)P(C)$

(c) $P(D|B,C)P(C|A)P(B|A)P(A)$

(d) $P(D|A,C)P(C|B)P(B|A)P(A)$

Q2. Given the following Bayesian network where all random variables are binary, calculate the joint probability $P(A,E)$ by performing belief propagation. Show the steps of how you get the results.

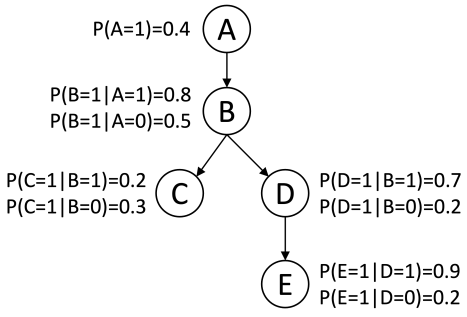

Problem 4: Reinforcement Learning (20 marks)

An agent lives in a 2×2 grid, and is using Q-learning to learn a policy, with the update rule:

$$Q(s,a) \leftarrow Q(s,a) + \eta \left(r + \left[\gamma \max_{a'} Q(s',a') \right] - Q(s,a) \right).$$

The state is the index of square that the agent locates. The agent will get reward $r = 10$ when it lands in State 3. There are no other rewards or penalties.

1	2
3 reward: 10	4

The agent has four possible actions, which are MoveEast, MoveWest, MoveNorth, and MoveSouth. The Q matrix is initialized to all zeros. The agent starts from Square 1 and makes the following actions: MoveSouth, MoveEast, MoveNorth, MoveWest. After each action, the Q-table is updated correspondingly. Assume that the learning rate $\eta = 0.5$ and the discount rate $\gamma = 0.9$. List out the nonzero entries of the Q-table at the end of this phase. Show your steps of making updates to the Q-table.