

UNIVERSITY OF BIRMINGHAM

School of Computer Science

Artificial Intelligence 2

Main Summer Examinations 2021

Artificial Intelligence 2

Exam paper

Question 1 (Probabilistic AI/ML)

As a data scientist in a telecommunication company, your task is to analyse a customer dataset to predict whether a customer will terminate his/her contract. The dataset consists of around 8000 customer records, each consisting of one binary dependent variable Y , indicating whether the customer terminates the contract ($Y = 1$) or not ($Y = 0$), and 19 independent variables, which include the customer's information, e.g., age, subscription plan, extra data plan, etc., and the consumer behaviour such as average numbers of calls and hours per week. Since your boss needs some actionable insights to retain customers, you decided to use interpretable machine learning methods. Design your interpretable machine learning method by answering the following questions:

- (a) You have implemented a feature selection algorithm based on mutual information to select the most informative features from the 19 independent variables. To validate the implementation of your mutual information calculation function, you use a small subset of the data to calculate mutual information manually. You select one independent variable 'subscription plan', denoted as S , which takes two values, $S \in \{1, 2\}$. Please use the following Probability Mass Function table

$p(S, Y)$	$S = 1$	$S = 2$
$Y = 0$	$\frac{2}{12}$	$\frac{5}{12}$
$Y = 1$	$\frac{2}{12}$	$\frac{3}{12}$

to calculate

- Entropies $H(S)$ and $H(Y)$
- Conditional entropies $H(S|Y)$ and $H(Y|S)$
- Joint entropy $H(S, Y)$
- Mutual information $I(S; Y)$

Show all your working. Discuss what mutual information means and whether this feature will be selected or not. **[6 marks]**

- (b) After applying your algorithm you selected two variables: 1) extra data plan E , which is a binary random variable that indicates whether the customer subscribes to the extra data plan ($E = 1$) or not ($E = 0$); and 2) averaged hours used per week H , which is a continuous random variable. You then built a logistic regression model

to classify customers into 'low risk' or 'high risk' of terminating the contract. The fitted model is

$$\log \left(\frac{p}{1-p} \right) = -0.77 + 0.23H - 1.18E$$

- Given a customer who has the extra data plan ($E = 1$) and spent on average 0.5 hours per week, calculate the odds and the probability the customer will terminate the contract ($Y = 1$). **[4 marks]**
- Using this fitted model, explain to your boss what actions should be taken to retain customers. **[10 marks]**

Question 2 (Deep Learning / Artificial Neural Networks)

- (a) Consider that we have a 64x64x5 dimensions medical image (i.e. 64x64 pixels and 5 channels) which we are processing through a convolutional layer of a deep convolutional neural network. Answer the following for this image:
- (i) Consider that the convolutional layer has 32 receptive fields each of size 7x7x5, without zero padding and a stride size of 1. What will be the output dimensions of the image produced from this convolutional layer? Justify your answer briefly with reasons through arguments and/or diagrams. **[6 marks]**
 - (ii) Consider that you pass the output obtained from the convolutional layer in (i) above to a channel-wise maximum pooling layer with a pool size 2x2 and a stride size of 2. Note that channel wise max pooling layer operates on each channel separately. What will be the output dimensions of the image produced from this max pooling layer? Justify your answer briefly with reasons through arguments and/or diagrams. **[4 marks]**
- (b) Consider that a cancer hospital intends to develop an AI solution for the automatic prediction of cancer from the magnetic resonance image (MRI) of a patient's brain. Formulate this problem as a deep machine learning problem and suggest how its solution can be developed. Your answer should include: problem formulation (e.g. what are the data & labelling needs), identification of relevant deep neural network that can be used for solution development (i.e. just the name of relevant network), the training/learning process of such network (i.e. how weights/parameters will be determined/updated) and the performance evaluation of the solution. Note: you are not required to explain a learning algorithm and neither are you required to provide mathematical formulae for the algorithm or performance evaluation.

[10 marks]