TE WHARE WĀNANGA O TE ŪPOKO O TE IKA A MĀUI
VICTORIA
UNIVERSITY OF WELLINGTON

EXAMINATIONS – 2020

TRIMESTER TWO

COMP 422

DATA MINING, NEURAL NETWORKS AND GENETIC PROGRAMMING

Time Allowed:

OPEN BOOK

Permitted materials: This is an open-book test. You can use any material for doing the test.

Instructions: You have 24 hours to work on the test paper.
There are a total of 60 marks on this test.
Attempt all questions.

Questions

	Marks
1. Neural Networks	[25]
2. Evolutionary Computation and Feature Manipulation	[18]
3. Feature Manipulation	[7]
4. Reinforcement Learning	[10]

1. **Neural Networks** **(25 marks)**

(a) **(2 marks)** State *two* neural network architectures for *supervised* learning and *two* neural network architectures for *unsupervised* learning.

(b) **(3 marks)** Convolutional neural network generates feature maps using the idea of convolution. It scans an image with a convolution mask. List *three types* of parameters that need to be set while generating a feature map.

(c) **(2 marks)** State a main difference between the *max pooling* and *convolution* filters.

(d) **(4 marks)** Briefly explain *two* advantages of transfer learning over standard supervised machine learning methods.

(e) **(4 marks)** James has developed a neural network for predicting the infection rate of a flu virus. He used the `monthOfYear` as a feature, and used one input node to represent it. The node can take integer values from 1 to 12, indicating the month of the year. However, the results showed that the neural network has poor generalisation especially between December of one year and January of the next year.

Suggest a new representation of `monthOfYear` to improve the neural network's generalisation. Draw a figure to show your suggestion if necessary.

(f) **(10 marks)** Chris wanted to design a convolutional neural network for image classification. The input image has 32x32 pixels. In the first hidden layer, he designed to generate 8 feature maps, each generated by using a 5x5 convolutional mask with shift of 2 pixels. There is no padding for the convolution.

(i) Calculate the number of rows and columns in each feature map. Show your working.

(ii) Calculate the total number of parameters (weights + biases) between the input layer and one feature map. Show your working.

(iii) Calculate the total number of *unique* parameters (weights + biases) between the input layer and one feature map. Show your working.

(iv) In the second hidden layer, Chris decided to generate 8 feature maps, each from 4 feature maps in the first hidden layer, using a 3x3 convolutional mask with shift of 1 pixel. There is no padding for the convolution. Calculate the total number of *unique* parameters between the first hidden layer and one feature map in the second hidden layer. Show your working.

2. **Evolutionary Computation** **(18 marks)**

(a) **(1 mark)** State *two* ways to increase the exploration ability of a genetic algorithm.

(b) **(1 mark)** Briefly describe the major differences between *genetic algorithm* and *genetic programming* regarding representation and the search process.

(c) **(1 mark)** Briefly describe the major differences between *genetic programming* and *particle swarm optimisation* in terms of representation and the search process.

(d) **(2 marks)** Particle Swarm Optimisation (PSO) typically uses the following formulae to do the search:

$$v_{id}^{t+1} = w * v_{id}^{t} + c_1 * r_1 * (p_{id} - x_{id}^{t}) + c_2 * r_2 * (p_{gd} - x_{id}^{t}) \quad (1)$$

$$x_{id}^{t+1} = x_{id}^{t} + v_{id}^{t+1} \quad (2)$$

Describe how the algorithm behaviour will change if c_1 is increased (or c_2 is decreased).

(e) **(4 marks)** Simon is developing a genetic algorithm for solving traveling salesman problem, which is to find the shortest cycle that visits each node in the graph exactly once.

i) He chooses the straightforward representation, which represents each individual as a permutation of all the cities (e.g., $[3, 2, 5, 4, 1]$ for a 5-node graph).
However, traditional genetic operators (e.g., single-point crossover operator and flip mutation) cannot guarantee the feasibility of offspring. Design a specialised crossover operator and a specialised mutation operator that can always generate feasible offspring. Clearly describe each of them. You can use examples to help your description.

ii) In order to directly use the traditional genetic operators, Simon changes the representation to a continuous vector, e.g., $[0.15, 2.11, -3.2, 1.67, 1.02]$ for a 5-node graph. Design a *decoding* scheme that can convert ANY continuous vector to a feasible solution to traveling salesman problem. Clearly describe the designed decoding scheme.

(f) **(4 marks)** Strongly-Typed Genetic Programming (STGP) can be used to solve the symbolic regression problem where the dataset contains variables with different types. In STGP, each terminal primitive is associated with a data type, and each function primitive contains the data type for each element in its parameter list, and the data type of its output.

Given the following dataset, in which $\{x_1, x_2, x_3, x_4\}$ are input variables and y is the output variable. x_1 and x_2 are boolean variables, while x_3 and x_4 are numerical.

Design the terminal set and function set for the STGP for this dataset. Clearly describe (1) each terminal and its associated data type, and (2) each function, the data type of each of its parameters, and the data type of its output.

x_1	x_2	x_3	x_4	y
true	true	2.5	1.5	4.75
true	false	3.1	2.2	7.41
false	false	2.2	1.1	2.0
false	false	4.3	4.3	1.0
...	...	...	...	...

(g) **(5 marks)** We have discussed Non-dominated Sorting Genetic Algorithm II (NSGA-II) for multi-objective optimization in our lecture. It selects the individuals based on (1) non-dominated sorting, and (2) crowding distance.

The figure below shows 12 individuals in the 2-dimensional objective space. Each node represents an individual. For example $A(2,6)$ indicates the individual A with $f1(A) = 2$ and $f2(A) = 6$.

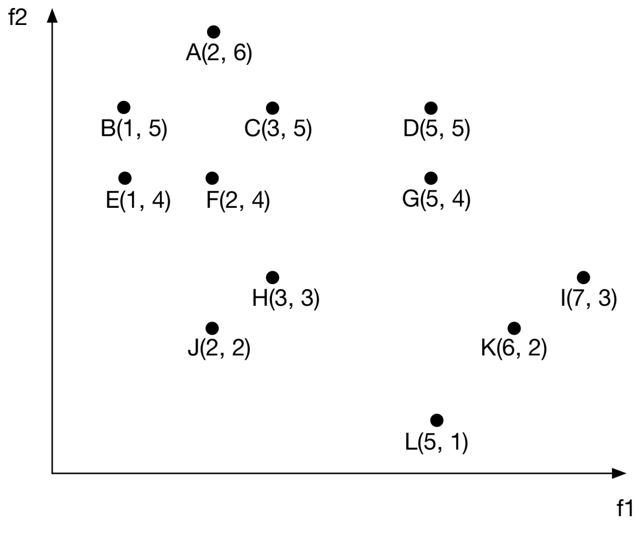

Assume it is a minimisation problem, which tries to minimise $f1$ and $f2$, describe the steps of NSGA-II for selecting 6 individuals in the figure for the next generation.

For each selected individual, give the rank calculated by the non-dominated sorting (start from rank 1) and the crowding distance (or explain why it is not required to calculate for this individual).

3. Feature Manipulation **(7 marks)**

(a) **(2 marks)** Given the following dataset containing 12 instances with a single feature

$$\{1.3, 1.4, 1.7, 2.5, 2.7, 2.8, 3.1, 3.2, 3.4, 3.6, 5.1, 5.3\},$$

calculate the *equal-width* and *equal-depth* discretisation results with 4 categories. Show your working.

(b) **(2 marks)** Suppose that a classification problem has three features A, B and C, and that using a classification algorithm, we obtain the following performances (classification accuracies) for all the possible subsets of features.

Subset	$\{\}$	$\{A\}$	$\{B\}$	$\{C\}$	$\{A,B\}$	$\{A,C\}$	$\{B,C\}$	$\{A,B,C\}$
Performance	0.5	0.62	0.6	0.66	0.88	0.78	0.82	0.79

i) Assuming the *sequential forward selection* algorithm is used, show the steps it takes to find a subset of features.

ii) Assuming the *sequential backward selection* algorithm is used, show the steps it takes to find a subset of features.

(c) **(1 mark)** Clearly describe a way to use *particle swarm optimisation for feature construction*.

(d) **(2 marks)** Clearly describe *two* ways to use *genetic programming for feature selection*.

4. **Reinforcement Learning and NeuroEvolution** **(10 marks)**

(a) **(1 mark)** Identify and briefly explain the major differences between model-based and model-free reinforcement learning.

(b) **(1 mark)** Briefly explain why model-free reinforcement learning algorithms are more popularly used in practice.

(c) **(2 marks)** State two main advantages of the OpenAI Evolution Strategy for coping with reinforcement learning problems.

(d) **(3 marks)** For any *stochastic* policy π, define the V-function and Q-function of policy π. Provide suitable mathematical equations to support the definition.

(e) **(3 marks)** Learning the Q-function of a *deterministic* policy π is often approached by using the following learning rule with respect to a single sampled state transition (s_t, s_{t+1}, r_t, a_t), where s_t and s_{t+1} refer to the states of the learning environment at time step t and $t+1$ respectively. Meanwhile, $a_t = \pi(s_t)$ stands for the action performed by the reinforcement learning agent at time step t and r_t denotes the corresponding immediate reward received from the learning environment.

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha \left(r_t + \gamma Q(s_{t+1}, \pi(s_{t+1})) - Q(s_t, a_t) \right)$$

In the above learning rule, α is the learning rate and γ is the discount factor. Now consider two consecutive state transition samples: (s_t, s_{t+1}, r_t, a_t) and $(s_{t+1}, s_{t+2}, r_{t+1}, a_{t+1})$ with $a_t = \pi(s_t)$ and $a_{t+1} = \pi(s_{t+1})$. Develop correspondingly a new rule for learning the Q-function of deterministic policy π. This new learning rule must use both of the two consecutive state transition samples.

* * * * * * * * * * * * * * *