ALGORITHMICS (HESS)

Written examination

February 3, 2024

Section	Number of questions	Number to be answered	Number of marks
A	20	20	20
B	15	15	89

Wi nøt trei a høliday in Sweden this yer?
See the løveli lakes
The wonderful telephøne system
And mani interesting furry animals
The Producers would like to thank The Forestry Commission Doune Admissions Ltd, Keir and Cowdor Estates, Stirling University, and the people of Doune for their help in the making of this film. The Characters and incidents portrayed and the names used are fictitious and any similarity to the names, characters, or history of any person is entirely accidental and unintentional. Signed RICHARD M. NIXON
Including the majestic møøse
A Møøse once bit my sister ...
No realli! She was Karving her initials on the møøse with the sharpened end of an interspace tøøthbrush given her by Svenge - her brother-in-law - an Oslo dentist and star of many Norwegian møvies: "The Høt Hands of an Oslo Dentist", "Fillings of Passion", "The Huge Mølars of Horst Nordfink".

Students are NOT permitted to bring mobile phones and/or any other unauthorised electronic devices into the examination room.

SECTION A – Multiple-choice questions

Question 1
One limitation of using SVMs to classify data is that they

A. Can only process linearly separable data correctly

B. Only allow data with a maximum of 3 features

C. Require very large training datasets compared to other machine learning models to make predictions

D. Only allow binary classification of data

Question 2
Which of the below MUST be true of computational complexity classes?

A. P $\subseteq$ NP

B. NP $\supset$ P

C. NP-Complete $\supset$ NP-Hard

D. NP-Hard $\subseteq$ P

Question 3
If a problem is proven to be undecidable then

A. No algorithm will give a correct output for any input

B. It is in computational complexity class NP-Complete

C. It is the Halting Problem

D. It must be a decision problem

Question 4
Which values of p, q, and z would cause the following statement to evaluate to false?
NOT((p AND q) OR (p AND z)) OR z

A. p = true, z = false, q = false

B. p = true, z = false, q = true

C. p = false, z = true, q = false

D. p = false, z = false, q = false

Question 5
King Arthur wishes to write a program that generates the number of ways to choose 3 coconuts from a pile of n coconuts, where all the coconuts are distinct. Suppose that he has written the most efficient algorithm possible, then the tightest bound on the big O of his algorithm could be

A. O(n!)

B. O(2^n)

C. O(n^2)

D. O(n^3)

Question 6
The Cart Master needs to visit every house in their village once and return back to the starting point. They wish to do this efficiently, without revisiting any houses. The problem that most closely resembles the Cart Master's problem is

A. finding an Euler circuit on an undirected graph

B. finding the minimal spanning tree of a graph

C. finding a Hamiltonian circuit for a graph

D. finding the transitive closure of a graph

Question 7
Tim is training his AI for voice recognition by recording people in a cafe without their knowledge. This main ethical issue of concern in this scenario is

A. Accountability
B. Transparency
C. Fairness
D. Bias

Question 8
Which of the below statements is true?

A. All directed acyclic graphs are connected.
B. A tree has $|V| = |E|$
C. A connected graph minimally has $|V| - 1 = |E|$
D. A connected graph must have $|E| = \frac{|V|(|V|-1)}{2}$

Question 9
Using the Turing test as a benchmark for AI capability would best support which response to Searle's Chinese room thought experiement?

A. Systems
B. Robots
C. Other minds
D. Neural network/ brain

Question 10
Consider the following graph. One problem with running Dijkstra's Algorithm on this graph is

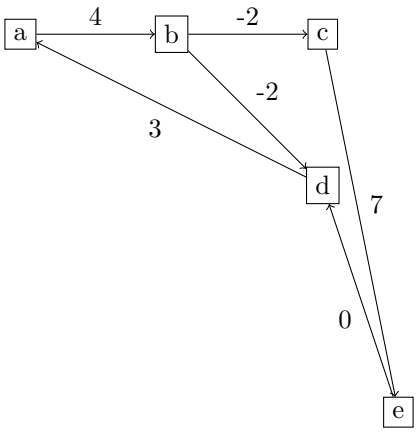

A. The existence of negative weight cycles
B. The existence of an edge with weight 0
C. The directed nature of the graph
D. The existence of negative weight edges

Question 11
One difference between the A* algorithm and the best-first search algorithm is that

A. Best-first search considers only the heuristic function when ranking nodes whereas A* also considers cost
B. A* uses a heuristic function but best-first search does not
C. The A* algorithm always gurantees an optimal solution unlike best-first search
D. Best-first search terminates after finding the goal whereas A* finds the shortest path from the goal to all other nodes

Question 12
The A* algorithm can guarantee the output of an optimal shortest path if and only if its heuristic

A. Predicts nodes' actual costs exactly

B. Underestimates the true cost of nodes and is non monotonic

C. Overestimates the true cost of nodes and is monotonic

D. Underestimates the true cost of nodes and is monotonic

Question 13
What does the algorithm ni return when n = 0 and a = 17?

```
Ni(a,n):
  if a = 1:
    return n
  else:
    if a is odd
      return ni(3a+1,n+1)
    else
      return ni(a/2, n+1)
```

A. 9

B. 10

C. 11

D. 12

Question 14
Consider the following list ADT signatures:

Append: elem ×_____ →list Update: _____ ×list →list

The blanks in the list ADT signatures are best filled in by.

A. Int, list

B. List, int

C. List, list

D. Int, int

Question 15
If a graph has multiple possible MSTs, then it must

A. have negative edge weights

B. have duplicate edge weights

C. be a complete graph

D. have a negative cycle

Question 16
If a graph G with verticies V and edges E is known to be bipartite, which of the following may be false?

A. G must have no cycles length 3

B. $|V|$ must be even

C. G cannot be complete if $|V| > 2$

D. If no neighbouring nodes can share a colour, G's vertices are able to be coloured with 3 colours

Question 17
Forward propogation in neural networks refers to

A. Minimising the loss function through adjusting edge weights

B. Creation of a new feature from existing features to ensure linear separability

C. Summation of features and bias within a perceptron

D. Transformation of data that gets passed from input to output

Question 18
Which of these criteria is true of all Turing Machines?

A. A Turing Machine must eventually halt

B. A Turing Machine may have a countably infinite set of states

C. A Turing Machine must have an infinite tape of cells.

D. A Turing Machine must have a transition function which maps every state-symbol pair to an action

Question 19
Zoot is a doctor who wishes to treat incoming patients according to how severe their conditions are, however, with limited staffing, she must also ensure that each patient is treated in a timely manner. She decides that no more than x patients who arrived later than some patient can be treated earlier than that patient. Which combination of ADTs most accurately models her strategy?

A. Dictionary of priority queues

B. Queue of priority queues

C. Stack of queues

D. Priority queue of stacks

Question 20
What is the big O time complexity of an algorithm with the recurrence relation $T(n) = 5T(\frac{n}{2}) + n^3$ if n>1 and O(1) otherwise?

A. n^3

B. $n^3 log(n)$

C. $n^{log_2(5)}$

D. n^2

SECTION B – Short answer questions

Question 1

(a) (3 points) Outline a minimal set of signature specifications for the array ADT

(b) (2 points) Outline a similarity and a difference between lists and arrays.

Question 2

(a) (2 points) Bedevere is cooking. He is making a very complex dish called "meat pie" where certain procedures – each consisting of multiple steps – must be completed before other procedures. For example, a pie crust must be created before filling can be put inside. Explain an ADT or combination of ADTs that can suitably model this problem.

(b) (3 points) Describe an algorithm with reference to your data model that will tell Bedevere the order in which to complete the recipe's steps.

Question 3

(a) (6 points) The Knights of the Round Table are trying to guess the passcode to be allowed to cross the Bridge of Death. The password is a 4-letter string which is known to consist of the letters a,b,c,d, and e. Assuming that they have an unlimited number of attempts at guessing, write an algorithm that will find the correct passcode. You may use the following function:

`check(n)`

– takes an input **list** n and will return true if each n[i] is equal to the ith letter of the passcode. I.e., if n represents the correct passcode.

(b) (2 points) Miraculously, the group uses your algorithm and happily makes it to the other side. Sir Galahad decides to implement facial recognition software on the bridge to prevent others from crossing, deeming the passcode too weak. Explain the advantages that neural networks have over traditional computational methods for problems like facial recognition.

Question 4

(a) (3 points) Using the definition "a tree is a connected undirected graph with no cycles", prove that the removal of any edge from a tree results in a disconnected graph.

Question 5

(a) (6 points) A "consecutive pair" occurs in a word when a letter is next to the same letter. For example, "spleen" has one consecutive pair "ee" and "balloon" has two consecutive pairs "ll" and "oo". Annabella is an avid enjoyer of the VCE subject English Language and wants to find the longest chain of *consecutive* consecutive pairs in some words for her essay. For example, "balloon" has a longest consecutive pair chain of two consecutive pairs in "lloo", whereas "hullabaloo" has a longest consecutive chain of only 1 as none of its consecutive pairs are consecutive. Write an algorithm in pseudocode that takes a list s of all the letters in a word in order (a la "cat" -> ["c","a","t"]) and outputs the number of consecutive consecutive pairs. You may assume that no word has three consecutive letters that are the same.

Question 6

Hilbert's program was a proposal for a complete formalisation of foundational mathematics. It was started by David Hilbert in the 19th century, and its goals included proofs for the "completeness", "consistency", and "decidability" of mathematics.

(a) (3 points) Outline what Hilbert's program aimed to achieve with regard to "decidability". How does this compare to "completeness"?

(b) (2 points) How do Turing Machines relate to the concept of decidability?

Question 7

(a) (4 points) VCAA wants to minimise the number of students who have exam clashes for the end of year exams. A student is said to have a clash if their exams overlap with each others' time slots; conversely, exams clash with one another if they share a time slot and share students. What ADT would be most appropriate to model exams and clashes? Assuming that there are only k exam slots, desrcibe an algorithm can be used to check if preventing all clashes if possible? And if not, how can you modify your approach such that the fewest number of students have exam clashes?

(b) (2 points) Evaluate the feasibility of your algorithm. DO NOT calculate an explicit time complexity.

Question 8

(a) (1 point) Dingo and Zoot are discussing different ways to find the transitive closure of a graph. Define transitive closure.

(b) (4 points) Dingo knows that the Floyd-Warshall transitive closure algorithm can find the transitive closure of a graph. Zoot responds, claiming that would be faster to simply run breadth-first search on every node. Is Zoot's solution more efficient than Dingo's? Is this true for every graph? Justify your answer.

Question 9

(a) (2 points) "Collect the Capital" is a single-player game played on a hexagonal board. The board may be hundreds of tiles wide. In the game, the player has k moves, where each move consists of moving to an adjacent hexagon. If the player moves into a hexagon with $c money, the player gains c points. The aim of the game is to collect as much money as possible in k moves. Outline a suitable ADT to model the game board with justification.

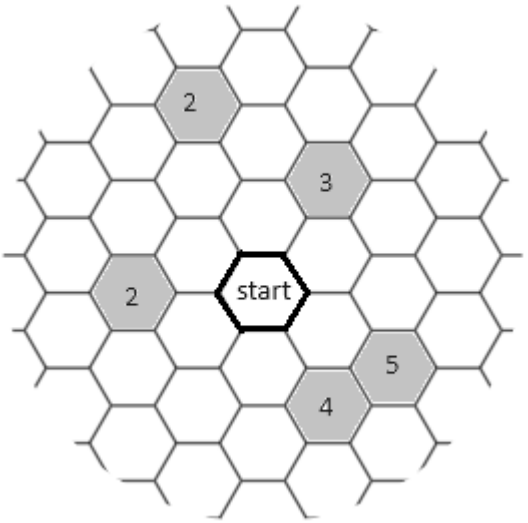

(b) (5 points) Write a brute force algorithm that determines the most optimal route to take to maximise monetary gain.

(c) (2 points) Explain how the practical solvability of this problem varies with the input size.

(d) (2 points) Simulated annealing is a design pattern that uses randomisation to generate a solution in feasible time. Explain the role of randomisation in finding close-to-optimal solutions for hard problems.

(e) (5 points) Write an algorithm in pseudocode that uses simulated annealing to find an acceptable solution, taking the following parameters as input:

- T – final temperature
- T_0 – initial temperature
- g – the goal profit
- d – the acceptable level difference between current and goal profit
- C – the cooling factor

Assume that the function p(T) exists and generates the probability of accepting a new solution based on the value of T – a higher T will give a higher probability.

Question 10

(a) (4 points) Explain how Floyd-Warshall's algorithm can be used to find the diameter of a weighted graph, justifying any modifications that need to be made.

Question 11

(a) (4 points) A tournament graph is a graph such that there is a directed edge connecting every pair of nodes. Prove by induction that a tournament graph of any size can be a directed acyclic graph.

Question 12

(a) (2 points) Some plants exhibit recursive patterns in growth. The Knights who say Ni! are a group which loves botany and their leader wants to model the growth of his favourite plant, parsley, using an algorithm. He knows that each time a main branch of parsley branches off into smaller branches, it branches into exactly 3 such branches. (This is true: take a look next time you go foraging in the woods!).

```
parsley(d):
  if d<1
    return
  else
    for i from 1 to 3
      make branch
      parsley(d-1)
```

Calculate with justification the big-O time complexity of the parsley algorithm if "make branch" runs in O(1) time.

(b) (2 points) The leader decides that this algorithm produces parsley that looks too uniform. He changes the algorithm to betterParsley displayed below, varying the growth of individual branches.

```
betterParsley(d):
  if d<1
    return
  else
    for i from 1 to 3
      make branch
      betterParsley(d-i)
```

Write a recursive relation to describe the time complexity of betterParsley. Is it more efficient to run for the same depth than the previous algorithm? Justify.

Question 13

(a) (2 points) *AI can now predict with over 90 percent accuracy a student's test score, identifying their weakness and strengths with about 10 minutes of interaction. We can predict what questions a student will get wrong before they even try to answer them. We can even predict when a student will get tired and disengage.*

(Marten Roorda, 2022)

The Minister of Education, Stewort Molner, is thinking about implementing AI in high schools to replace a portion of standardised testing. Currently, students' academic results are standardised to be comparable through undertaking these tests, but concerns have been raised about the validity of testing in assessing understanding – as well as issues such socioeconomic status which is correlated with test results. To warn him of potential errors made by AI, provide a description of the concepts of model overfitting and underfitting for the minister.

(b) (3 points) The minister is philosophizing about the "true intelligence" of AI. He has been reading about Searle's Chinese Room thought experiment and says: "A big difference between our brains and AI is the sheer number of neurons that make up human brains. I wonder if a complex enough system could truly model a brain and have intelligence." Which of the standard responses to the thought experiment is most similar to the minister's statement? Make an argument for or against this response.

(c) (2 points) Explain the difference between forward propogation and backpropogation in neural networks.

(d) (1 point) What is the role of the activation function for a perceptron?

(e) (10 points) The Ministry of Education has behavioural data and cross-subject academic data for students from the last 20 years. It can only retain personal information data (such as gender and disability) from the last 5 years,

so drastically fewer datasets are available. If an AI testing system is implemented, it will judge students based upon virtual conversations, data about classroom behaviour, and streamed testing. In an ideal world, ethical decision-making ould be integral to the core tenets of government, so it is important to the ministry that the process of judging students based on AI is transparent, unbiased, and upholds privacy.
Consider the options of using an SVM-based AI, using a neural-network based AI, and retaining the current system of human-based evaluation and testing. Provide a recommendation as to which system for judging learning outcomes would be most suitable by evaluating:

- all three options and their relationship to the ministry's ethical values
- choices for training data for AI systems
- the expected quality of decision making