

ADSSE Extra re-exam 23 October 2018

Question 1

- a) True
- b) False
- c) True
- d) False
- e) True
- f) False
- g) True
- h) True
- i) True
- j) True

Question 2

- a) $T(n) = \Theta(n^2)$
- b) $T(n) = \Theta(n^3 \log_2 n)$
- c) $T(n) = \Theta((\sqrt{n})^5)$
- d) $T(n) = \Theta(\sqrt{n} \log_2 n)$
- e) $T(n) = \Theta(n^4)$

Question 3

- a) A
- b) C
- c) B
- d) D
- e) B

Question 4

- a) 5 7 9
- b) 3 blocks per grid, 1 thread per block
- c) CPU memory
- d) device memory
- e) 27-29

Question 5

- a) $\Theta(n^2)$. Occurs when input array is completely sorted.
- b) $\Theta(n \log_2 n)$
- c) Linear – occurs when input array is already sorted.
- d) $\Theta(n^2)$ – occurs when input array is reverse sorted.
- e) $O(n)$

Question 6

- a) A machine learning algorithm is trained by using a labelled sample of inputs called the training set. The parameters of the model are adjusted until it outputs the correct label for the training set in a sufficiently high proportion of cases. The trained model is then used to predict the label of previously unseen inputs.
- b) In k-fold cross-validation, the data is randomly partitioned into k equally sized, non-intersecting subsets. Each of k steps, the model is trained on k-1 of these parts and then tested on the remaining part. The average of the accuracies achieved by the model over the k folds is reported. All observations are used for both training and validation.

- c) In an ensemble method, a set of two or more independent classifiers are used to predict the class of an input. The predictions of the different models in the ensemble are combined to produce a single prediction. This can be done using either bagging or boosting.
- d) Bagging is where the predictions of multiple classifiers in an ensemble are averaged (majority vote). m new training sets are generated by random sampling and m models are fitted using the m training sets.
- e) Boosting is where the prediction of a classifier in an ensemble is weighted more heavily if it is more accurate.

Question 7

- a) Because a simple perceptron sums its inputs and outputs a 1 if the sum of the inputs is above a threshold. An XOR gate has to output a 0 when there is a 1 on both inputs and a 1 when there is a 1 on only one of the inputs. However if there is a 1 on both inputs, the sum is always going to be higher than if there is a 1 on only one of the inputs, so a perceptron can implement OR but not XOR.
- b) The network is given a training sample, x , and the output of the network, y , is observed. y is compared with y' , the desired output for input x , and a measure is computed for the difference between y and y' (that is, the error). The weights and biases in the network are modified to reduce the error. The process is repeated for different training inputs x until either a certain number of training epochs has been completed or the error is reduced to a satisfactory level.
- c) Mutation is an operator that performs random alteration in a candidate solution and produces one new individual. Crossover swaps two homologous segments in the parent genotypes to produce two new child individuals – it is the process of merging information from two parent genotypes into one offspring.

Question 8

- a)
 - i. User
 - ii. User
 - iii. Service provider
 - iv. Service provider
- b) Multi-tenancy is a principle in software architecture where a single instance of the software serves multiple tenants or customers.
- c) Vertical scaling means increasing the size and power of the machines used; horizontal scaling means using more machines in a grid architecture. Horizontal scaling in the data tier is usually achieved through sharding, which means splitting a database across two or more databases to handle higher query or data volumes.
- d) A CDN is a network of distributed servers in multiple data centers in different locations around the globe. Users can then obtain data from the nearest server in the CDN which allows the data to be delivered more quickly.
- e) Polyglot persistence is using multiple data storage technologies, chosen based upon the way data is being used by applications. Apache Spark is an example of software that enables polyglot persistence.

Question 9

- a) 3
[3, 4]
[[1, 2, 3, 4]]
[11, 12]
- b) Two of the following: Euclidean, city-block (manhattan), cosine
- c) A collision occurs when two or more keys have the same hash value and are therefore sent to the same slot in the hash table.

- d) Each slot in the hash table points at the head of a linked list containing the keys that have the hash value corresponding to that slot in the table. Each time a new key is entered in the table, it is placed at the head of the linked list hanging off the slot for the key's hash value. A collision therefore results in such a linked list containing more than one element.
- e) $O(1)$, constant time.

Question 10

Should be something like this:

```
9 def min_max(d):
10     i = 0
11     low = []
12     hi = []
13     while i < len(d)-1:
14         if d[i] < d[i+1]:
15             low.append(d[i])
16             hi.append(d[i+1])
17         else:
18             low.append(d[i+1])
19             hi.append(d[i])
20         i+=2
21     if i < len(d):
22         low.append(d[i])
23         hi.append(d[i])
24     min_val = low[0]
25     j = 1
26     while j < len(low):
27         if low[j] < min_val:
28             min_val = low[j]
29         j+=1
30     max_val = hi[0]
31     k = 1
32     while k < len(hi):
33         if hi[k] > max_val:
34             max_val = hi[k]
35         k+=1
36     print("Max = ", max_val)
37     print("Min = ", min_val)
```