# ADSSE Re-exam 2018 Answers

## Question 1

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

[1 mark each]

## Question 2

a) $\Theta(n\, log_2\, n)$
b) $\Theta(\log_2 \mathrm{n})$
c) $\Theta(n^2)$
d) $\Theta\left(n^{log_2 7}\right)$
e) $\Theta(n\, log_2\, n)$

[2 marks each]

## Question 3

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

[2 marks each]

## Question 4

a)
```
 0  1  2  3
 4  5  6  7
 8  9 10 11
12 13 14 15
```

b) 4 blocks per grid, 4 threads per block
c) CPU memory
d) Device memory
e) 17

[2 marks for each part]

## Question 5

a) INSERT, SEARCH (TEST FOR MEMBERSHIP), DELETE

b) A collision occurs when two keys hash to the same slot in the hash table – that is, they have the same hash value.
c) Collisions can be handled by *chaining* (also called closed hashing) or open hashing. In chaining, each slot in the hash table has a pointer to a linked list. When a new input key is hashed to the slot, a new list node is created and inserted at the head of the hash slot's linked list. (We didn't cover open addressing.)
d) Constant time, O(1).
e) The load factor of a hash table is the ratio n/m where n is the number of keys in the data and m is the number of slots in the hash table. Simple uniform hashing occurs when there is an equal probability for any given key hashing to any slot in the hash table.

[2 marks each]

## Question 6

a) True
b) False
c) True
d) False
e) False

[2 marks each]

## Question 7

a) X = Input layer, Y = Hidden layer, Z = Output layer. [3 marks, 1 mark each]
b)
   i. relu or rectified linear unit [1 mark]

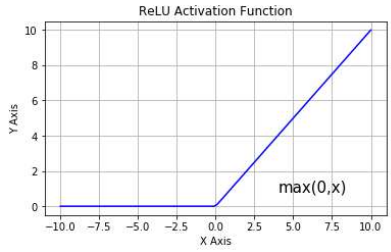


[1 mark]

   ii. A dense layer is fully connected to the next layer in the network (every node in a dense layer has a connection to every node in the next layer). [1 mark]
   iii. Uses a binary cross-entropy loss function. [1 mark] Could have used any of the loss functions listed on the following page: https://keras.io/losses/ [1 mark]
   iv. 1 [1 mark]
   v. 20 [1 mark]

## Question 8

a) False
b) False
c) True
d) True
e) True

[2 marks for each part]

## Question 9

a) The three models are as follows [1 mark for each]

i. Infrastructure as a Service (IaaS) provides a capability for users to provision processing, storage, and network resources on demand. The customers deploy and run their own applications on these resources. Examples include Amazon EC2 cloud computing, Amazon S3 storage services.
ii. In Platform as a Service (PaaS) the service provider makes certain core components such as databases, queues, workflow engines, email, and so on, available as services to the customer. The customer then leverages these components for building their own applications. The service provider ensures high service levels, and is responsible for scalability, high availability, and so on, for these components. An example is Amazon DynamoDB.
iii. In Software as a Service (SaaS), typically, third-party providers using a subscription model provide end user applications to their customers. The customers may have some administrative capability at the application level, for example, to create and manage their users. Examples include machine learning, computer vision (e.g., Amazon Rekognition).

b) An Amazon Machine Image (AMI) provides the information required to launch an instance, which is virtual server in the cloud.
Or
An AMI is a special type of virtual appliance that is used to create a virtual machine within the Amazon Elastic Computing Cloud (EC2). [1 mark]
c) Amazon Rekognition is a cloud-based system that uses deep learning for classifying and labelling images and videos. [2 marks]
d) An AWS Region is a geographic location where AWS provides multiple physically separated and isolated Availability Zones that are connected with low latency, high throughput and highly redundant networking. [2 marks]
e) AWS S3 is object storage built to store and retrieve any amount of data from anywhere on the Internet. It is a simple storage service that offers a durable, available, affordable and scalable data storage infrastructure. [2 marks]

## Question 10

a) Eventual consistency is a consistency model used in distributed computing to achieve high availability that informally guarantees that, if no new updates are made to a given data item, eventually all accesses to that item will return the last updated value. It is fairly common in the cloud world. After a data update, if your application can tolerate a few seconds delay before the update is reflected across all replicas of the data then eventual consistency can lead to better scalability and performance.
b) Sharding is a type of database partitioning that separates very large databases into smaller, faster, more easily managed parts called data shards.
c) AWS Identity and Access Management (IAM) enables you to manage access to AWS services and resources securely.
d) A keypair is public-key cryptography pair to encrypt and decrypt login information for your Amazon EC2 instance.
e) Serverless computing allows you to build and run applications and services without thinking about servers. With serverless computing, your application still runs on servers, but all the server management is done by AWS. At the core of serverless computing is AWS Lambda, which lets you run your code without provisioning or managing servers.

[2 marks for each correct part]