



**Free Questions for 1Z0-1122-23 by certsinside**

**Shared by Vaughan on 22-07-2024**

**For More Free Questions and Preparation Resources**

**Check the Links on Last Page**

# Question 1

---

**Question Type:** MultipleChoice

---

What is the primary purpose of Convolutional Neural Networks (CNNs)?

## Options:

---

- A- Processing sequential data
- B- Generating Images
- C- Creating music compositions
- D- Detecting patterns in images

## Answer:

---

D

## Explanation:

---

Convolutional Neural Networks (CNNs) are a type of deep learning algorithm that is particularly well-suited for image recognition and processing tasks. They are made up of multiple layers, including convolutional layers, pooling layers, and fully connected layers. The convolutional layer is the core building block of a CNN, and it is where the majority of computation occurs. It requires a few components,

which are input data, a filter, and a feature map. The filter is a small matrix of weights that slides over the input data and performs element-wise multiplication and summation, resulting in a feature map that represents the activation of a certain feature in the input. By applying multiple filters, the CNN can detect different patterns in the image, such as edges, shapes, colors, textures, etc. The pooling layer is used to reduce the spatial dimensionality of the feature maps, while preserving the most important information. The fully connected layer is the final layer of a CNN, and it is where the classification or regression task is performed based on the extracted features. CNNs can learn to detect complex patterns in images by adjusting their weights during training using backpropagation and gradient descent algorithms. Reference::Convolutional neural network - Wikipedia,What are Convolutional Neural Networks? | IBM,Convolutional Neural Network (CNN) in Machine Learning

## Question 2

---

**Question Type:** MultipleChoice

---

What is the purpose of fine-tuning Large Language Models?

**Options:**

---

- A-** To reduce the number of parameters in the model
- B-** To Increase the complexity of the model architecture

**C-** To specialize the model's capabilities for specific tasks

**D-** To prevent the model from overfitting

**Answer:**

---

C

**Explanation:**

---

Fine-tuning is the process of updating the model parameters on a new task and dataset, using a pre-trained large language model as the starting point. Fine-tuning allows the model to adapt to the specific context and domain of the new task, and improve its performance and accuracy. Fine-tuning can be used to customize the model's capabilities for specific tasks such as text classification, named entity recognition, and machine translation<sup>82</sup>. Fine-tuning is also known as transfer learning or task-based learning. Reference: A Complete Guide to Fine Tuning Large Language Models, Finetuning Large Language Models - DeepLearning.AI

## Question 3

---

**Question Type:** MultipleChoice

---

How is "Prompt Engineering" different from "Fine-tuning" in the context of Large Language Models (LLMs)?

### Options:

---

- A- Customizes the model architecture
- B- Trains a model from scratch
- C- Guides the model's response using predefined prompts
- D- Involves post-processing model outputs and optimizing hyper parameters

### Answer:

---

C

### Explanation:

---

Prompt engineering is the art of designing natural language instructions or queries that can elicit the desired response from a large language model. Prompt engineering does not modify the model parameters or architecture, but rather relies on the model's existing knowledge and capabilities. Prompt engineering can be used to perform various tasks such as text generation, sentiment analysis, and code completion, by providing the model with the appropriate context, format, and constraints<sup>67</sup>. Prompt engineering is also known as zero-shot learning or query-based learning. Reference: [2211.01910] Large Language Models Are Human-Level Prompt Engineers](<https://arxiv.org/abs/2211.01910>), A developer's guide to prompt engineering and LLMs - The GitHub Blog

## Question 4

---

**Question Type: MultipleChoice**

---

What is "in-context learning" in the realm of large Language Models (LLMs)?

**Options:**

---

- A- Teaching a model through zero-shot learning
- B- Training a model on a diverse range of tasks
- C- Modifying the behavior of a pretrained LLM permanently
- D- Providing a few examples of a target task via the input prompt

**Answer:**

---

D

**Explanation:**

---

In-context learning is a technique that leverages the ability of large language models to learn from a few input-output examples provided in the input prompt. By conditioning on these examples, the model can infer the task and the format of the desired output, and generate a suitable response. In-context learning does not require any additional training or fine-tuning of the model, and can be used for various tasks such as text summarization, question answering, text generation, and more<sup>45</sup>. In-context learning is also known as few-shot learning or prompt-based learning. Reference: [2307.12375] In-Context Learning in Large Language Models Learns Label ...](<https://arxiv.org/abs/2307.12375>), [2307.07164] Learning to Retrieve In-Context Examples for Large Language

## Question 5

---

**Question Type:** MultipleChoice

---

Which NVIDIA GPU is offered by Oracle Cloud Infrastructure?

**Options:**

---

**A-** P200

**B-** T4

**C-** A100

**D-** K80

**Answer:**

---

C

## **Explanation:**

---

Oracle Cloud Infrastructure offers NVIDIA A100 Tensor Core GPUs as one of the GPU options for its compute instances. The NVIDIA A100 GPU is a powerful and versatile GPU that can accelerate a wide range of AI and HPC workloads. The A100 GPU delivers up to 20x higher performance than the previous generation V100 GPU and supports features such as multi-instance GPU, automatic mixed precision, and sparsity acceleration<sup>12</sup>. The OCI Compute bare-metal BM.GPU4.8 instance offers eight 40GB NVIDIA A100 GPUs linked via high-speed NVIDIA NVLink direct GPU-to-GPU interconnects<sup>3</sup>. This instance is ideal for training large language models, computer vision models, and other complex AI tasks. Reference: Accelerated Computing and Oracle Cloud Infrastructure (OCI) - NVIDIA, Oracle Cloud Infrastructure Offers New NVIDIA GPU-Accelerated Compute ..., GPU, Virtual Machines and Bare Metal | Oracle

## **Question 6**

---

**Question Type:** MultipleChoice

---

What is the advantage of using Oracle Cloud Infrastructure Supercluster for AI workloads?

## **Options:**

---

**A-** It offers seamless integration with social media platforms.



- B-** It provides a cost-effective solution for simple AI tasks.
- C-** It delivers exceptional performance and scalability for complex AI tasks.
- D-** It is ideal for tasks such as text-to-speech conversion.

**Answer:**

---

C

**Explanation:**

---

Oracle Cloud Infrastructure Supercluster is a cloud service that provides ultrafast cluster networking, HPC storage, and OCI Compute bare metal instances. OCI Supercluster is ideal for training generative AI, including conversational applications and diffusion models, as it can deploy up to tens of thousands of NVIDIA GPUs per cluster for much greater scalability than similar offerings from other providers. OCI Supercluster also reduces the time needed to train AI models with simple Ethernet network architecture that provides ultrahigh performance at massive scale. Additionally, OCI Supercluster offers cost savings and access to AI subject matter experts<sup>56</sup>. Reference: OCI Supercluster and AI Infrastructure | Oracle, Oracle Delivers More Choices for AI Infrastructure and General-Purpose ...

## Question 7

---

**Question Type:** MultipleChoice

---

You are the lead developer of a Deep Learning research team, and you are tasked with improving the training speed of your deep neural networks. To accelerate the training process, you decide to leverage specialized hardware.

Which hardware component is commonly used in Deep Learning to accelerate model training?

**Options:**

---

- A- Solid-State Drive (SSD)
- B- Graphics Processing Unit (GPU)
- C- Random Access Memory (RAM)
- D- Central Processing Unit (CPU)

**Answer:**

---

B

**Explanation:**

---

A graphics processing unit (GPU) is a specialized hardware component that can perform parallel computations on large amounts of data. GPUs are widely used in deep learning to accelerate the training of deep neural networks, as they can execute many matrix operations and tensor operations simultaneously. GPUs can significantly reduce the training time and improve the performance of deep learning models compared to using CPUs alone<sup>678</sup>. Reference: [Hardware Recommendations for Machine Learning / AI](#), [New hardware offers faster computation for artificial intelligence ...](#), [The Best Hardware for Machine Learning - ReHack](#), [Hardware for Deep Learning](#)

## Question 8

---

**Question Type:** MultipleChoice

---

What is the difference between Large Language Models (LLMs) and traditional machine learning models?

**Options:**

---

- A- LLMs require labeled output for training.
- B- LLMs have a limited number of parameters compared to other models.
- C- LLMs are specifically designed for natural language processing and understanding.
- D- LLMs focus on image recognition tasks.

**Answer:**

---

C

## Explanation:

---

Large language models (LLMs) are a class of deep learning models that can recognize and generate natural language, among other tasks. LLMs are trained on huge sets of text data, learning grammar, semantics, and context. LLMs use the Transformer architecture, which relies on self-attention to process and understand the input and output sequences. LLMs can perform various natural language processing and understanding tasks based on the input provided, such as text summarization, question answering, text generation, and more<sup>34</sup>. Traditional machine learning models, on the other hand, are usually trained with specific statistical algorithms that deliver pre-defined outcomes. They often require labeled data and feature engineering, and they are not as flexible and adaptable as LLMs<sup>5</sup>. Reference: What are LLMs, and how are they used in generative AI?, An Introduction to LLM Ops: Operationalizing and Managing Large Language Models using Azure ML, An Introduction to Large Language Models (LLMs): How It Got ... - Labellerr

## Question 9

---

**Question Type:** MultipleChoice

---

What is the purpose of Attention Mechanism in Transformer architecture?

### Options:

---

**A-** Convert tokens into numerical forms (vectors) that the model can understand.

- B-** Break down a sentence into smaller pieces called tokens.
- C-** Apply a specific function to each word individually.
- D-** Weigh the importance of different words within a sequence and understand the context.

**Answer:**

---

D

**Explanation:**

---

The attention mechanism in the Transformer architecture is a technique that allows the model to focus on the most relevant parts of the input and output sequences. It computes a weighted sum of the input or output embeddings, where the weights indicate how much each word contributes to the representation of the current word. The attention mechanism helps the model capture the long-range dependencies and the semantic relationships between words in a sequence<sup>12</sup>. Reference: [The Transformer Attention Mechanism - MachineLearningMastery.com](#), [Attention Mechanism in the Transformers Model - Baeldung](#)

**To Get Premium Files for 1Z0-1122-23 Visit**

**<https://www.p2pexams.com/products/1z0-1122-23>**

**For More Free Questions Visit**

**<https://www.p2pexams.com/oracle/pdf/1z0-1122-23>**

