Free Questions for 1Z0-1122-23

Shared by Nieves on 04-10-2023

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

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

Options:

A- Teaching a mode! 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 more45. 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 Models](https://arxiv.org/abs/2307.07164)

Question 2

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 more34. 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 LLMs5.Reference:What are LLMs, and how are they used in generative AI?,An Introduction to LLMOps: Operationalizing and Managing Large Language Models using Azure ML,An Introduction to Large Language Models (LLMs): How It Got ... - Labellerr

Question 3

Question Type: MultipleChoice

What is the primary function of Oracle Cloud Infrastructure Speech service?

Options:

A- Converting text into images
B- Analyzing sentiment n text
C- Transcribing spoken language into written text
D- Recognizing objects in images

Answer:

C

Explanation:

Oracle Cloud Infrastructure Speech is an AI service that applies automatic speech recognition (ASR) technology to transform audio-based content into text. Developers can easily make API calls to integrate Speech's pretrained models into their applications. Speech can be used for accurate, text-normalized, time-stamped transcription via the console and REST APIs as well as command-line interfaces or SDKs. You can also use Speech in an OCI Data Science notebook session.With Speech, you can filter profanities, get confidence scores for both single words and complete transcriptions, and more1.Reference:Speech AI Service that Uses ASR | OCI Speech - Oracle

Question 4

Question Type: MultipleChoice

Which capability is supported by Oracle Cloud Infrastructure Language service?

Options:

A- Analyzing text to extract structured information like sentiment or entities
B- Detecting objects and scenes in Images
C- Translating speech into text
D- Converting text into images

Answer:

A

Explanation:

Oracle Cloud Infrastructure Language service is a cloud-based AI service for performing sophisticated text analysis at scale. It provides various capabilities to process unstructured text and extract structured information like sentiment or entities using natural language processing techniques. Some of the capabilities supported by Oracle Cloud Infrastructure Language service are:

Language Detection: Detects languages based on the provided text, and includes a confidence score.

Text Classification: Identifies the document category and subcategory that the text belongs to.

Named Entity Recognition: Identifies common entities, people, places, locations, email, and so

on.

Key Phrase Extraction: Extracts an important set of phrases from a block of text.

Sentiment Analysis: Identifies aspects from the provided text and classifies each into positive, negative, or neutral polarity.

Text Translation: Translates text into the language of your choice.

Personal Identifiable Information: Identifies, classifies, and de-identifies private information in unstructured textReference::Language Overview - Oracle,AI Text Analysis at Scale | Oracle

Question 5

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 alone678.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

Inference: How to Choose the Best One for ...

Question 6

Question Type: MultipleChoice

What is the primary purpose of reinforcement learning?

Options:

A- Finding relationships within data sets
B- Identifying patterns in data
C- Making predictions from labeled data
D- Learning from outcomes to make decisions

Answer:

D

Explanation:

Reinforcement learning is a type of machine learning that is based on learning from outcomes to make decisions. Reinforcement learning algorithms learn from their own actions and experiences in an environment, rather than from labeled data or explicit feedback. The goal of reinforcement learning is to find an optimal policy that maximizes a cumulative reward over time. A policy is a rule that determines what action to take in each state of the environment. A reward is a feedback signal that indicates how good or bad an action was for achieving a desired objective. Reinforcement learning involves a trial-and-error process of exploring different actions and observing their consequences, and then updating the policy accordingly. Some of the challenges and components of reinforcement learning are:

Exploration vs exploitation: Balancing between trying new actions that might lead to higher rewards in the future (exploration) and choosing known actions that yield immediate rewards (exploitation).

Markov decision process (MDP): A mathematical framework for modeling sequential decision making problems under uncertainty, where the outcomes depend only on the current state and action, not on the previous ones.

Value function: A function that estimates the expected long-term return of each state or state-action pair, based on the current policy.

Q-learning: A popular reinforcement learning algorithm that learns a value function called Q-function, which represents the quality of taking a certain action in a certain state.

Deep reinforcement learning: A branch of reinforcement learning that combines deep neural networks with reinforcement learning algorithms to handle complex and high-dimensional problems, such as playing video games or controlling robots.Reference::Reinforcement learning - Wikipedia,What is Reinforcement Learning? -- Overview of How it Works - Synopsys

Question 7

Question Type: MultipleChoice

Which capability is supported by the Oracle Cloud Infrastructure Vision service?

Options:

A- Detecting and classifying objects in images
B- Generating realistic Images from text
C- Analyzing historical data for unusual patterns
D- Detecting and preventing fraud in financial transactions

Answer:

A

Explanation:

Oracle Cloud Infrastructure Vision is a serverless, multi-tenant service, accessible using the Console, or over REST APIs. You can upload images to detect and classify objects in them. If you have lots of images, you can process them in batch using asynchronous API endpoints. Vision's features are thematically split between Document AI for document-centric images, and Image Analysis for object and scene-based images.Image Analysis supports both pretrained and custom models for object detection and image classification3.Reference:Vision - Oracle

Question 8

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 translation82. 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 9

Question Type: MultipleChoice

Which is an application of Generative Adversarial Networks (GANs) in the context of Generative AI?

Options:

A- Creation of realistic images that resemble training data
B- Prediction of continuous values from Input data
C- Generation of labeled outputs for training
D- Classification of data points into categories

Answer:

A

Explanation:

Generative Adversarial Networks (GANs) are a type of AI model that can generate realistic images that resemble training data. The architecture of a GAN consists of two separate neural networks that are pitted against each other in a game-like scenario. The first network, known as the generator network, tries to create fake data that looks real. The second network, known as the discriminator network, tries to distinguish between real and fake data. The generator network learns from the feedback of the discriminator network and tries to fool it by improving the quality of the fake data. The discriminator network also learns from the feedback of the generator network and tries to improve its accuracy.The process continues until the generator network produces data that is indistinguishable from the real data4.GANs can be used to create realistic images of faces, animals, landscapes, and more5.Reference:Generative models - OpenAI,Artificial Intelligence Explained: What Are Generative Adversarial ...

P2P exams