Give Push to your Success with Oracle Cloud 1z0-1122-23 Exam Questions [Q17-Q35

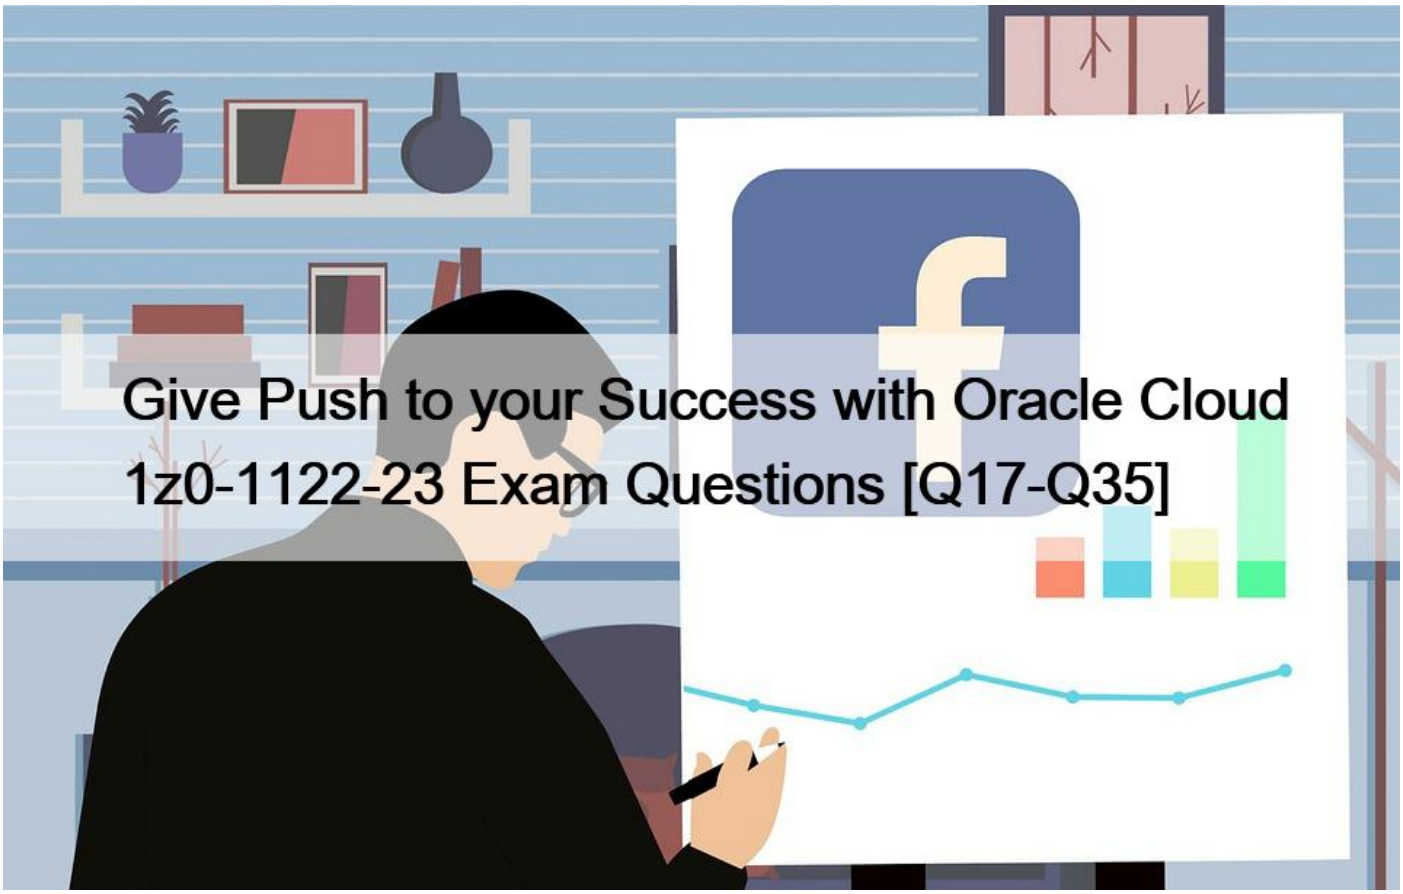

Give Push to your Success with Oracle Cloud 1z0-1122-23 Exam Questions
1z0-1122-23 100% Guarantee Download 1z0-1122-23 Exam PDF Q&A

NEW QUESTION 17

What is the difference between Large Language Models (LLMs) and traditional machine learning models?
* LLMs require labeled output for training.
* LLMs have a limited number of parameters compared to other models.
* LLMs are specifically designed for natural language processing and understanding.
* LLMs focus on image recognition tasks.
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

NEW QUESTION 18

What is the difference between classification and regression in Supervised Machine Learning?
* Classification assigns data points to categories, whereas regression predicts continuous values.
* Classification and regression both predict continuous values.
* Classification predicts continuous values, whereas regression assigns data points to categories.
* Classification and regression both assign data points to categories.

Classification and regression are two subtypes of supervised learning in machine learning. The main difference between them is the type of output variable they deal with. Classification assigns data points to discrete categories based on some criteria or rules. For example, classifying emails into spam or not spam based on their content is a classification problem because the output variable is binary (spam or not spam). Regression predicts continuous values for data points based on their input features. For example, predicting house prices based on their size, location, amenities, etc., is a regression problem because the output variable is continuous (house price). Classification and regression use different types of algorithms and metrics to evaluate their performance. Reference: : Oracle Cloud Infrastructure AI – Machine Learning Concepts, Classification vs Regression in Machine Learning | by …

NEW QUESTION 19

Which AI domain is associated with tasks such as recognizing forces in images and classifying objects?
* Computer Vision
* Anomaly Detection
* Speech Processing
* Natural Language Processing

Computer Vision is an AI domain that is associated with tasks such as recognizing faces in images and classifying objects. Computer vision is a field of artificial intelligence that enables computers and systems to derive meaningful information from digital images, videos, and other visual inputs, and to take actions or make recommendations based on that information. Computer vision works by applying machine learning and deep learning models to visual data, such as pixels, colors, shapes, textures, etc., and extracting features and patterns that can be used for various purposes. Some of the common techniques and applications of computer vision are:

Face recognition: Identifying or verifying the identity of a person based on their facial features.

Object detection: Locating and labeling objects of interest in an image or a video.

Object recognition: Classifying objects into predefined categories, such as animals, vehicles, fruits, etc.

Scene understanding: Analyzing the context and semantics of a visual scene, such as the location, time, weather, activity, etc.

Image segmentation: Partitioning an image into multiple regions that share similar characteristics, such as color, texture, shape, etc.

Image enhancement: Improving the quality or appearance of an image by applying filters, transformations, or corrections.

Image generation: Creating realistic or stylized images from scratch or based on some input data, such as sketches, captions, or attributes. Reference: : What is Computer Vision? | IBM, Computer vision – Wikipedia

NEW QUESTION 20

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

* Detecting and classifying objects in images
* Generating realistic Images from text
* Analyzing historical data for unusual patterns
* Detecting and preventing fraud in financial transactions

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

NEW QUESTION 21

Which AI task involves audio generation from text?

* Text to speech
* Audio recording
* Speech recognition
* Text summarization

Text to speech (TTS) is an AI task that involves audio generation from text. TTS is a technology that converts text into spoken audio using natural sounding voices. TTS can read aloud any text data, such as PDFs, websites, books, emails, etc., and provide an auditory format for accessing written content. TTS can be helpful for anyone who needs to listen to text data for various reasons, such as accessibility, convenience, multitasking, learning, entertainment, etc. TTS uses different techniques and models to generate speech from text data, such as:

Concatenative synthesis: Combining pre-recorded segments of human speech based on the phonetic units of the text.

Parametric synthesis: Generating speech signals from acoustic parameters derived from the text using statistical models.

Neural synthesis: Using deep neural networks to learn the mapping between text and speech features and produce high-quality speech signals.

Expressive synthesis: Adding emotions or styles to the speech output to make it more natural and engaging. Reference: : Text-to-Speech AI: Lifelike Speech Synthesis | Google Cloud, Text-to-speech synthesis – Wikipedia

NEW QUESTION 22

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

* Creation of realistic images that resemble training data
* Prediction of continuous values from Input data
* Generation of labeled outputs for training
* Classification of data points into categories

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 …

NEW QUESTION 23

Which capability is supported by Oracle Cloud Infrastructure Language service?
* Analyzing text to extract structured information like sentiment or entities
* Detecting objects and scenes in Images
* Translating speech into text
* Converting text into images

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 text Reference: : Language Overview – Oracle, AI Text Analysis at Scale | Oracle

NEW QUESTION 24

How is Generative AI different from other AI approaches?
* Generative AI understands underlying data and creates new examples.
* Generative AI focuses on decision-making and optimization.
* Generative AI generates labeled outputs for training.
* Generative AI is used exclusively for text-based applications.

Generative AI is a branch of artificial intelligence that focuses on creating new content or data based on the patterns and structure of existing data. Unlike other AI approaches that aim to recognize, classify, or predict data, generative AI aims to generate data that is realistic, diverse, and novel. Generative AI can produce various types of content, such as images, text, audio, video, software code, product designs, and more. Generative AI uses different techniques and models to learn from data and generate new examples, such as generative adversarial networks (GANs), variational autoencoders (VAEs), diffusion models, and foundation models. Generative AI has many applications across different domains and industries, such as art, entertainment, education, healthcare, engineering, marketing, and more. Reference: : Oracle Cloud Infrastructure AI – Generative AI, Generative artificial intelligence – Wikipedia

NEW QUESTION 25

How is “Prompt Engineering” different from “Fine-tuning” in the context of Large Language Models (LLMs)?
* Customizes the model architecture
* Trains a model from scratch
* Guides the model’s response using predefined prompts
* Involves post-processing model outputs and optimizing hyper parameters

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 constraints67. 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

NEW QUESTION 26

What is the primary purpose of reinforcement learning?
* Finding relationships within data sets
* Identifying patterns in data
* Making predictions from labeled data
* Learning from outcomes to make decisions

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

NEW QUESTION 27

Which NVIDIA GPU is offered by Oracle Cloud Infrastructure?
* P200
* T4
* A100
* K80

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 acceleration12. 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 interconnects3. 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

NEW QUESTION 28

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?
* Solid-State Drive (SSD)
* Graphics Processing Unit (GPU)
* Random Access Memory (RAM)
* Central Processing Unit (CPU)

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 …