2024 Latest 1z0-1122-23 dumps Exam Material with 30 Questions [Q16-Q33

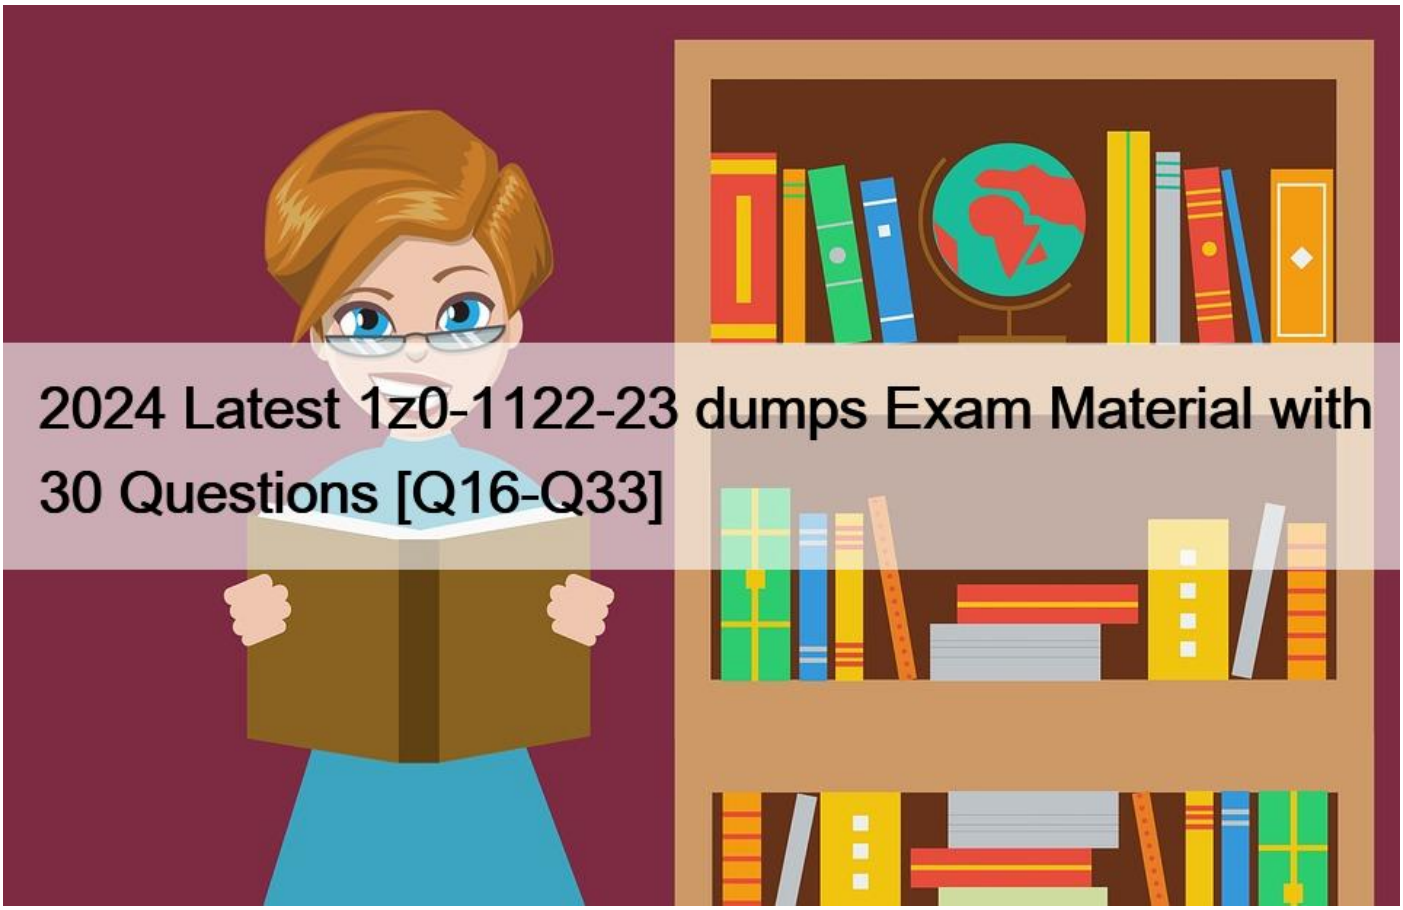

2024 Latest 1z0-1122-23 dumps Exam Material with 30 Questions Oracle 1z0-1122-23 Questions and Answers Guarantee you Oass the Test Easily NO.16 What is the advantage of using Oracle Cloud Infrastructure Supercluster for AI workloads?

* It offers seamless integration with social media platforms.
* It provides a cost-effective solution for simple AI tasks.
* It delivers exceptional performance and scalability for complex AI tasks.
* It is ideal for tasks such as text-to-speech conversion.

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 experts56. Reference: OCI Supercluster and AI Infrastructure | Oracle, Oracle Delivers More Choices for AI Infrastructure and General-Purpose …

NO.17 What is the primary purpose of Convolutional Neural Networks (CNNs)?

* Processing sequential data
* Generating Images
* Creating music compositions
* Detecting patterns in images

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

NO.18 You are working on a project for a healthcare organization that wants to develop a system to predict the severity of patients’ illnesses upon admission to a hospital. The goal is to classify patients into three categories – Low Risk, Moderate Risk, and High Risk – based on their medical history and vital signs.

Which type of supervised learning algorithm is required in this scenario?
* Clustering
* Regression
* Binary Classification
* Multi-Class Classification

Multi-class classification is a type of supervised learning algorithm that is required in this scenario because the output variable has more than two classes. Multi-class classification is the problem of classifying instances into one of three or more classes. For example, classifying patients into low risk, moderate risk, or high risk based on their medical history and vital signs is a multi-class classification problem because each patient can only belong to one of these three classes. Multi-class classification can be solved by using various algorithms, such as decision trees, random forests, support vector machines (SVMs), k-nearest neighbors (k-NN), naive Bayes, logistic regression, neural networks, etc. Some of these algorithms can naturally handle multi-class problems, while others need to be adapted by using strategies such as one-vs-one or one-vs-rest. Reference: : Multiclass classification – Wikipedia, Multiclass Classification- Explained in Machine Learning

NO.19 What is the primary function of Oracle Cloud Infrastructure Speech service?
* Converting text into images
* Analyzing sentiment n text
* Transcribing spoken language into written text
* Recognizing objects in images

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

NO.20 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

NO.21 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

NO.22 Which Deep Learning model is well-suited for processing sequential data, such as sentences?
* Generative Adversarial Network (GAN)
* Variational Autoencoder (VAE)
* Recurrent Neural Network (RNN)
* Convolutional Neural Network (CNN)

Recurrent Neural Networks (RNNs) are a type of deep learning algorithm that can process sequential data, such as sentences, speech, or time series. They are composed of recurrent units that have a loop that allows them to store information from previous inputs and pass it to the next inputs. This way, they can capture the temporal dependencies and context within a sequence. RNNs can be used for various natural language processing tasks, such as text generation, machine translation, sentiment analysis, speech recognition, etc. However, RNNs also suffer from some limitations, such as vanishing or exploding gradients, difficulty in modeling long-term dependencies, and high computational cost. Therefore, some variants and extensions of RNNs have been proposed to overcome these challenges, such as Long Short-Term Memory (LSTM), Gated Recurrent Unit (GRU), Bidirectional RNN (BiRNN), Attention Mechanism, etc. Reference: : [Recurrent neural network – Wikipedia], [What are Recurrent Neural Networks? | IBM], [Recurrent Neural Network (RNN) in Machine Learning]

NO.23 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

NO.24 Which type of machine learning is used to understand relationships within data and is not focused on making predictions or classifications?
* Active learning
* Unsupervised learning
* Reinforcement learning
* Supervised learning

Unsupervised learning is a type of machine learning that is used to understand relationships within data and is not focused on making predictions or classifications. Unsupervised learning algorithms work with unlabeled data, which means the data does not have predefined categories or outcomes. The goal of unsupervised learning is to discover hidden patterns, structures, or features in the data that can provide valuable insights or reduce complexity. Some of the common techniques and applications of unsupervised learning are:

Clustering: Grouping similar data points together based on their attributes or distances. For example, clustering can be used to segment customers based on their preferences, behavior, or demographics.

Dimensionality reduction: Reducing the number of variables or features in a dataset while preserving the essential information. For example, dimensionality reduction can be used to compress images, remove noise, or visualize high-dimensional data in lower dimensions.

Anomaly detection: Identifying outliers or abnormal data points that deviate from the normal distribution or behavior of the data. For example, anomaly detection can be used to detect fraud, network intrusion, or system failure.

Association rule mining: Finding rules that describe how variables or items are related or co-occur in a dataset. For example, association rule mining can be used to discover frequent itemsets in market basket analysis or recommend products based on purchase history. Reference: : Unsupervised learning – Wikipedia, What is Unsupervised Learning? | IBM

NO.25 What is the primary goal of machine learning?
* Enabling computers to learn and improve from experience
* Explicitly programming computers
* Creating algorithms to solve complex problems
* Improving computer hardware

Machine learning is a branch of artificial intelligence that enables computers to learn from data and experience without being explicitly programmed. Machine learning algorithms can adapt to new data and situations and improve their performance over time2. Reference: Artificial Intelligence (AI) | Oracle

NO.26 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

NO.27 Which AI domain is associated with tasks such as identifying the sentiment of text and translating text between languages?
* Natural Language Processing
* Speech Processing
* Anomaly Detection
* Computer Vision

Natural Language Processing (NLP) is an AI domain that is associated with tasks such as identifying the sentiment of text and translating text between languages. NLP is an interdisciplinary field that combines computer science, linguistics, and artificial intelligence to enable computers to process and understand natural language data, such as text or speech. NLP involves various techniques and applications, such as:

Text analysis: Extracting meaningful information from text data, such as keywords, entities, topics, sentiments, emotions, etc.

Text generation: Producing natural language text from structured or unstructured data, such as summaries, captions, headlines, stories, etc.

Machine translation: Translating text or speech from one language to another automatically and accurately.

Question answering: Retrieving relevant answers to natural language questions from a knowledge base or a document collection.

Speech recognition: Converting speech signals into text or commands.

Speech synthesis: Converting text into speech signals with natural sounding voices.

Natural language understanding: Interpreting the meaning and intent of natural language inputs and generating appropriate responses.

Natural language generation: Creating natural language outputs that are coherent, fluent, and relevant to the context. Reference: : What is Natural Language Processing? | IBM, Natural language processing – Wikipedia