

Accurate Hot Selling 1z0-1122-23 Exam Dumps 2023 Newly Released [Q15-Q31]



Accurate Hot Selling 1z0-1122-23 Exam Dumps 2023 Newly Released [Q15-Q31]

Accurate Hot Selling 1z0-1122-23 Exam Dumps 2023 Newly Released
Get 100% Authentic Oracle 1z0-1122-23 Dumps with Correct Answers

NEW QUESTION 15

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

NEW QUESTION 16

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 17

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 18

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

NEW QUESTION 19

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

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

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>)

NEW QUESTION 20

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 …

NEW QUESTION 21

What role do tokens play in Large Language Models (LLMs)?

- * They represent the numerical values of model parameters.
- * They are used to define the architecture of the model’s neural network.
- * They are individual units into which a piece of text is divided during processing by the model.
- * They determine the size of the model’s memory.

Tokens are the basic units of text representation in large language models. They can be words, subwords, characters, or symbols. Tokens are used to encode the input text into numerical vectors that can be processed by the model’s neural network. Tokens

also determine the vocabulary size and the maximum sequence length of the model³. Reference: Oracle Cloud Infrastructure 2023 AI Foundations Associate | Oracle University

NEW QUESTION 22

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 23

How does Oracle Cloud Infrastructure Anomaly Detection service contribute to fraud detection?

- * By identifying abnormal patterns in data
- * By analyzing text sentiment
- * By generating spoken language from text
- * By transcribing spoken language

Oracle Cloud Infrastructure Anomaly Detection is an AI service that provides real-time and batch anomaly detection for univariate and multivariate time series data. Through a simple user interface, organizations can create and train models to detect anomalies and identify unusual behavior, changes in trends, outliers, and more. Anomaly Detection can contribute to fraud detection by analyzing data from various sources, such as transactions, logs, sensors, or customer behavior, and alerting users when suspicious or fraudulent activities are detected². Reference: Anomaly Detection | Oracle

NEW QUESTION 24

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

- * To reduce the number of parameters in the model
- * To Increase the complexity of the model architecture
- * To specialize the model's capabilities for specific tasks
- * To prevent the model from overfitting

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

NEW QUESTION 25

Which AI domain is associated with tasks such as recognizing faces 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 26

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

Dumps of 1z0-1122-23 Cover all the requirements of the Real Exam:

<https://www.dumpsmaterials.com/1z0-1122-23-real-torrent.html>