

NCA-GENL Study Braindumps Make You Pass NCA-GENL Exam Fluently - Exam4PDF

NCA-GENL

QUESTION: 13

You are working on a regression problem to predict house prices based on several features, including the number of bedrooms, square footage, and neighborhood quality (categorical). Which combination of Python packages and methods should you use to prepare the dataset for a Linear Regression model?

Option A : Using pandas for data handling, LabelEncoder from scikit-learn for the categorical feature, and StandardScaler for scaling numerical features.

Option B : Using Keras for one-hot encoding of categorical features and TensorFlow for scaling numerical data.

Option C : Using scikit-learn's MinMaxScaler for all features, including categorical data, before applying the model.

Option D : Using spaCy to lemmatize the categorical features and NumPy to normalize the numerical features.

Correct Answer: A

QUESTION: 14

You are experimenting with two different generative AI models for summarizing legal documents. To determine which model performs better, you decide to compare them using statistical performance metrics. Which of the following metrics and methods should you prioritize for a meaningful comparison? (Select two)

Option A : Assess the models' performance using the Log-Likelihood metric.

Option B : Use the Area Under the Curve (AUC) to compare the models' performance.

Option C : Calculate the Mean Absolute Error (MAE) for the summaries generated by each model.

Option D : Measure the models' performance using the Total Variation Distance (TVD).

Option E : Evaluate the models using ROUGE (Recall-Oriented Understudy for Gisting Evaluation) scores.

Correct Answer: A,E

QUESTION: 15

You are developing a generative AI model that needs to generate high-quality images from textual descriptions in real-time. Which two of the following approaches will best optimize the performance and quality of your model given the hardware constraints of a GPU with limited memory? (Select two)

Option A : Implement mixed precision training

Option B : Increase the batch size to the maximum the GPU can handle

Option C : Implement data parallelism across multiple GPUs

Option D : Use gradient checkpointing

Option E : Use a larger learning rate to converge faster

Correct Answer: A,D

Link In comment

P.S. Free & New NCA-GENL dumps are available on Google Drive shared by Exam4PDF: https://drive.google.com/open?id=1qngXjP2LPEcnxUp_YnDwUk5xTDK5gDuX

Our online test engine and the windows software of the NCA-GENL study materials can evaluate your exercises of the virtual exam and practice exam intelligently. Our calculation system of the NCA-GENL study materials is designed subtly. Our evaluation process is absolutely correct. We are strictly in accordance with the detailed grading rules of the real exam. The point of every question is set separately. Once you submit your exercises of the NCA-GENL Study Materials, the calculation system will soon start to work.

NVIDIA NCA-GENL Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"> Software Development: This section of the exam measures the skills of Machine Learning Developers and covers writing efficient, modular, and scalable code for AI applications. It includes software engineering principles, version control, testing, and documentation practices relevant to LLM-based development.
Topic 2	<ul style="list-style-type: none"> LLM Integration and Deployment: This section of the exam measures skills of AI Platform Engineers and covers connecting LLMs with applications or services through APIs, and deploying them securely and efficiently at scale. It also includes considerations for latency, cost, monitoring, and updates in production environments.
Topic 3	<ul style="list-style-type: none"> Data Preprocessing and Feature Engineering: This section of the exam measures the skills of Data Engineers and covers preparing raw data into usable formats for model training or fine-tuning. It includes cleaning, normalizing, tokenizing, and feature extraction methods essential to building robust LLM pipelines.
Topic 4	<ul style="list-style-type: none"> Data Analysis and Visualization: This section of the exam measures the skills of Data Scientists and covers interpreting, cleaning, and presenting data through visual storytelling. It emphasizes how to use visualization to extract insights and evaluate model behavior, performance, or training data patterns.
Topic 5	<ul style="list-style-type: none"> Fundamentals of Machine Learning and Neural Networks: This section of the exam measures the skills of AI Researchers and covers the foundational principles behind machine learning and neural networks, focusing on how these concepts underpin the development of large language models (LLMs). It ensures the learner understands the basic structure and learning mechanisms involved in training generative AI systems.
Topic 6	<ul style="list-style-type: none"> Experiment Design
Topic 7	<ul style="list-style-type: none"> Alignment: This section of the exam measures the skills of AI Policy Engineers and covers techniques to align LLM outputs with human intentions and values. It includes safety mechanisms, ethical safeguards, and tuning strategies to reduce harmful, biased, or inaccurate results from models.
Topic 8	<ul style="list-style-type: none"> Prompt Engineering: This section of the exam measures the skills of Prompt Designers and covers how to craft effective prompts that guide LLMs to produce desired outputs. It focuses on prompt strategies, formatting, and iterative refinement techniques used in both development and real-world applications of LLMs.
Topic 9	<ul style="list-style-type: none"> Experimentation: This section of the exam measures the skills of ML Engineers and covers how to conduct structured experiments with LLMs. It involves setting up test cases, tracking performance metrics, and making informed decisions based on experimental outcomes.:
Topic 10	<ul style="list-style-type: none"> Python Libraries for LLMs: This section of the exam measures skills of LLM Developers and covers using Python tools and frameworks like Hugging Face Transformers, LangChain, and PyTorch to build, fine-tune, and deploy large language models. It focuses on practical implementation and ecosystem familiarity.

>> NCA-GENL Valid Test Camp <<

Updated NCA-GENL Dumps | NCA-GENL Exam Paper Pdf

Obtaining this NCA-GENL certificate is not an easy task, especially for those who are busy every day. However, if you use our NCA-GENL exam torrent, we will provide you with a comprehensive service to overcome your difficulties and effectively improve your ability. If you can take the time to learn about our NCA-GENL Quiz prep, I believe you will be interested in our NCA-GENL exam questions. Our NCA-GENL learning materials are practically tested, choosing our NCA-GENL exam guide, you will get unexpected surprise.

NVIDIA Generative AI LLMs Sample Questions (Q96-Q101):

NEW QUESTION # 96

What is the main difference between forward diffusion and reverse diffusion in diffusion models of Generative AI?

- A. Forward diffusion uses bottom-up processing, while reverse diffusion uses top-down processing to generate samples from noise vectors.
- B. Forward diffusion focuses on generating a sample from a given noise vector, while reverse diffusion reverses the process by estimating the latent space representation of a given sample.
- C. Forward diffusion uses feed-forward networks, while reverse diffusion uses recurrent networks.
- D. Forward diffusion focuses on progressively injecting noise into data, while reverse diffusion focuses on generating new samples from the given noise vectors.

Answer: D

Explanation:

Diffusion models, a class of generative AI models, operate in two phases: forward diffusion and reverse diffusion. According to NVIDIA's documentation on generative AI (e.g., in the context of NVIDIA's work on generative models), forward diffusion progressively injects noise into a data sample (e.g., an image or text embedding) over multiple steps, transforming it into a noise distribution. Reverse diffusion, conversely, starts with a noise vector and iteratively denoises it to generate a new sample that resembles the training data distribution. This process is central to models like DDPM (Denoising Diffusion Probabilistic Models). Option A is incorrect, as forward diffusion adds noise, not generates samples. Option B is false, as diffusion models typically use convolutional or transformer-based architectures, not recurrent networks. Option C is misleading, as diffusion does not align with bottom-up/top-down processing paradigms.

References:

NVIDIA Generative AI Documentation: <https://www.nvidia.com/en-us/ai-data-science/generative-ai/> Ho, J., et al. (2020). "Denoising Diffusion Probabilistic Models."

NEW QUESTION # 97

Which of the following claims is correct about quantization in the context of Deep Learning? (Pick the 2 correct responses)

- A. Helps reduce memory requirements and achieve better cache utilization.
- B. It only involves reducing the number of bits of the parameters.
- C. Quantization might help in saving power and reducing heat production.
- D. It consists of removing a quantity of weights whose values are zero.
- E. It leads to a substantial loss of model accuracy.

Answer: A,C

Explanation:

Quantization in deep learning involves reducing the precision of model weights and activations (e.g., from 32-bit floating-point to 8-bit integers) to optimize performance. According to NVIDIA's documentation on model optimization and deployment (e.g., TensorRT and Triton Inference Server), quantization offers several benefits:

* Option A: Quantization reduces power consumption and heat production by lowering the computational intensity of operations, making it ideal for edge devices.

References:

NVIDIA TensorRT Documentation: <https://docs.nvidia.com/deeplearning/tensorrt/developer-guide/index.html> NVIDIA Triton Inference Server Documentation: <https://docs.nvidia.com/deeplearning/triton-inference-server/user-guide/docs/index.html>

NEW QUESTION # 98

"Hallucinations" is a term coined to describe when LLM models produce what?

- A. Correct sounding results that are wrong.
- B. Grammatically incorrect or broken outputs.
- C. Outputs are only similar to the input data.
- D. Images from a prompt description.

Answer: A

Explanation:

In the context of LLMs, "hallucinations" refer to outputs that sound plausible and correct but are factually incorrect or fabricated, as emphasized in NVIDIA's Generative AI and LLMs course. This occurs when models generate responses based on patterns in training data without grounding in factual knowledge, leading to misleading or invented information. Option A is incorrect, as

hallucinations are not about similarity to input data but about factual inaccuracies. Option B is wrong, as hallucinations typically refer to text, not image generation. Option D is inaccurate, as hallucinations are grammatically coherent but factually wrong. The course states: "Hallucinations in LLMs occur when models produce correct-sounding but factually incorrect outputs, posing challenges for ensuring trustworthy AI." References: NVIDIA Building Transformer-Based Natural Language Processing Applications course; NVIDIA Introduction to Transformer-Based Natural Language Processing.

NEW QUESTION # 99

Which of the following optimizations are provided by TensorRT? (Choose two.)

- **A. Layer Fusion**
- B. Data augmentation
- C. Residual connections
- **D. Multi-Stream Execution**
- E. Variable learning rate

Answer: A,D

Explanation:

NVIDIA TensorRT provides optimizations to enhance the performance of deep learning models during inference, as detailed in NVIDIA's Generative AI and LLMs course. Two key optimizations are multi-stream execution and layer fusion. Multi-stream execution allows parallel processing of multiple input streams on the GPU, improving throughput for concurrent inference tasks. Layer fusion combines multiple layers of a neural network (e.g., convolution and activation) into a single operation, reducing memory access and computation time. Option A, data augmentation, is incorrect, as it is a preprocessing technique, not a TensorRT optimization. Option B, variable learning rate, is a training technique, not relevant to inference. Option E, residual connections, is a model architecture feature, not a TensorRT optimization. The course states:

"TensorRT optimizes inference through techniques like layer fusion, which combines operations to reduce overhead, and multi-stream execution, which enables parallel processing for higher throughput." References: NVIDIA Building Transformer-Based Natural Language Processing Applications course; NVIDIA Introduction to Transformer-Based Natural Language Processing.

NEW QUESTION # 100

What is Retrieval Augmented Generation (RAG)?

- A. RAG is a method for manipulating and generating text-based data using Transformer-based LLMs.
- **B. RAG is a methodology that combines an information retrieval component with a response generator.**
- C. RAG is an architecture used to optimize the output of an LLM by retraining the model with domain- specific data.
- D. RAG is a technique used to fine-tune pre-trained LLMs for improved performance.

Answer: B

Explanation:

Retrieval-Augmented Generation (RAG) is a methodology that enhances the performance of large language models (LLMs) by integrating an information retrieval component with a generative model. As described in the seminal paper by Lewis et al. (2020), RAG retrieves relevant documents from an external knowledge base (e.g., using dense vector representations) and uses them to inform the generative process, enabling more accurate and contextually relevant responses. NVIDIA's documentation on generative AI workflows, particularly in the context of NeMo and Triton Inference Server, highlights RAG as a technique to improve LLM outputs by grounding them in external data, especially for tasks requiring factual accuracy or domain- specific knowledge. Option A is incorrect because RAG does not involve retraining the model but rather augments it with retrieved data. Option C is too vague and does not capture the retrieval aspect, while Option D refers to fine-tuning, which is a separate process.

References:

Lewis, P., et al. (2020). "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks." NVIDIA NeMo Documentation: <https://docs.nvidia.com/deeplearning/nemo/user-guide/docs/en/stable/nlp/intro.html>

NEW QUESTION # 101

.....

It is evident to all that the NCA-GENL test torrent from our company has a high quality all the time. A lot of people who have bought our products can agree that our NCA-GENL test questions are very useful for them to get the certification. There have been

99 percent people used our NCA-GENL exam prep that have passed their exam and get the certification, more importantly, there are signs that this number is increasing slightly. It means that our NCA-GENL Test Questions are very useful for all people to achieve their dreams, and the high quality of our NCA-GENL exam prep is one insurmountable problem.

Updated NCA-GENL Dumps: <https://www.exam4pdf.com/NCA-GENL-dumps-torrent.html>

- NCA-GENL Questions □ NCA-GENL Latest Exam Cram □ NCA-GENL Certification Materials □ Open [www.examdisscuss.com] and search for 「 NCA-GENL 」 to download exam materials for free □ Practice NCA-GENL Exam Online
- New NCA-GENL Test Labs □ Valid Braindumps NCA-GENL Free □ NCA-GENL Customizable Exam Mode □ > www.pdfvce.com □ is best website to obtain▷ NCA-GENL ◁ for free download □ Latest NCA-GENL Test Prep
- NCA-GENL Actual Dumps □ Trustworthy NCA-GENL Pdf □ Valid Braindumps NCA-GENL Free □ Go to website { www.passcollection.com } open and search for ☀ NCA-GENL □☀□ to download for free □ NCA-GENL Reliable Cram Materials
- Free PDF NCA-GENL - The Best NVIDIA Generative AI LLMs Valid Test Camp □ Easily obtain free download of ➡ NCA-GENL □□□ by searching on □ www.pdfvce.com □ □ Trustworthy NCA-GENL Pdf
- Pass Your NVIDIA NCA-GENL Exam with Excellent NCA-GENL Valid Test Camp Certainly □ Search for □ NCA-GENL □ and download exam materials for free through 《 www.examsreviews.com 》 □ Test NCA-GENL Passing Score
- NCA-GENL valid test torrent - NCA-GENL reliable test vce - NCA-GENL training pdf dumps □ Search on □ www.pdfvce.com □ for 「 NCA-GENL 」 to obtain exam materials for free download □ New NCA-GENL Test Vce Free
- Free PDF NCA-GENL - The Best NVIDIA Generative AI LLMs Valid Test Camp □ Search on▶ www.pass4leader.com ◀ for 《 NCA-GENL 》 to obtain exam materials for free download □ NCA-GENL Certification Materials
- Free PDF NCA-GENL - The Best NVIDIA Generative AI LLMs Valid Test Camp □ Download （ NCA-GENL ） for free by simply entering ✓ www.pdfvce.com □✓□ website □ NCA-GENL Customizable Exam Mode
- NCA-GENL Customizable Exam Mode □ NCA-GENL Certification Materials □ NCA-GENL Customizable Exam Mode □ Immediately open 【 www.exam4pdf.com 】 and search for （ NCA-GENL ） to obtain a free download □ □ NCA-GENL Latest Mock Exam
- Test NCA-GENL Passing Score □ NCA-GENL Questions □ Practice NCA-GENL Exam Online □ Copy URL 《 www.pdfvce.com 》 open and search for [NCA-GENL] to download for free □ NCA-GENL Latest Exam Cram
- NCA-GENL Actual Dumps □ Practice NCA-GENL Exam Online □ Practice NCA-GENL Exam Online □ Enter ☀ www.examcollectionpass.com □☀□ and search for ▶ NCA-GENL □ to download for free □ NCA-GENL Exam Forum
- www.stes.tyc.edu.tw, myportal.utt.edu.tw, myportal.utt.edu.tw, myportal.utt.edu.tw, myportal.utt.edu.tw, myportal.utt.edu.tw, myportal.utt.edu.tw, myportal.utt.edu.tw, myportal.utt.edu.tw, myportal.utt.edu.tw, tahike9295.blogolenta.com, pruebas.alquiniaregenerativa.com, academy.gti.com.ng, www.stes.tyc.edu.tw, adewde.free-blogz.com, lingopediamagazin.com, 123.59.83.120:8080, a.zhhxq.cn, Disposable vapes

P.S. Free 2025 NVIDIA NCA-GENL dumps are available on Google Drive shared by Exam4PDF: https://drive.google.com/open?id=1qngXjP2LPEcnxUp_YnDwUk5xTDK5gDuX