

# NCA-GENL Exam Details | NCA-GENL Dump Check

## 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

BTW, DOWNLOAD part of Dumpkiller NCA-GENL dumps from Cloud Storage: <https://drive.google.com/open?id=1QM7-TC2LmR18BzeojMMIM4DgMhq0MKrC>

We know deeply that a reliable NCA-GENL exam material is our company's foothold in this competitive market. High accuracy and high quality are the most important things we always looking for. Compared with the other products in the market, our NCA-GENL latest questions grasp of the core knowledge and key point of the real exam, the targeted and efficient NVIDIA Generative AI LLMs study training dumps guarantee our candidates to pass the test easily. Our NCA-GENL Latest Questions is one of the most wonderful reviewing NVIDIA Generative AI LLMs study training dumps in our industry, so choose us, and together we will make a brighter future.

## NVIDIA NCA-GENL Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"> <li>• 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.</li> </ul>
Topic 2	<ul style="list-style-type: none"> <li>• 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.:</li> </ul>
Topic 3	<ul style="list-style-type: none"> <li>• 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.</li> </ul>
Topic 4	<ul style="list-style-type: none"> <li>• Experiment Design</li> </ul>
Topic 5	<ul style="list-style-type: none"> <li>• 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.</li> </ul>
Topic 6	<ul style="list-style-type: none"> <li>• 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.</li> </ul>
Topic 7	<ul style="list-style-type: none"> <li>• 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.</li> </ul>
Topic 8	<ul style="list-style-type: none"> <li>• 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.</li> </ul>
Topic 9	<ul style="list-style-type: none"> <li>• 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.</li> </ul>
Topic 10	<ul style="list-style-type: none"> <li>• This section of the exam measures skills of AI Product Developers and covers how to strategically plan experiments that validate hypotheses, compare model variations, or test model responses. It focuses on structure, controls, and variables in experimentation.</li> </ul>

>> NCA-GENL Exam Details <<

## NCA-GENL Dump Check & NCA-GENL Latest Braindumps Ppt

The NVIDIA Generative AI LLMs NCA-GENL exam dumps are top-rated and real NVIDIA Generative AI LLMs NCA-GENL practice questions that will enable you to pass the final NVIDIA Generative AI LLMs NCA-GENL exam easily. With the NVIDIA Generative AI LLMs Exam Questions you can make this task simple, quick, and instant. Using the NVIDIA Generative AI LLMs NCA-GENL can help you success in your exam. Dumpkiller offers reliable guide files and reliable exam guide materials for 365 days free updates.

### NVIDIA Generative AI LLMs Sample Questions (Q80-Q85):

#### NEW QUESTION # 80

Which Python library is specifically designed for working with large language models (LLMs)?

- A. HuggingFace Transformers
- B. NumPy
- C. Pandas
- D. Scikit-learn

**Answer: A**

Explanation:

The HuggingFace Transformers library is specifically designed for working with large language models (LLMs), providing tools for model training, fine-tuning, and inference with transformer-based architectures (e.g., BERT, GPT, T5). NVIDIA's NeMo documentation often references HuggingFace Transformers for NLP tasks, as it supports integration with NVIDIA GPUs and frameworks like PyTorch for optimized performance.

Option A (NumPy) is for numerical computations, not LLMs. Option B (Pandas) is for data manipulation, not model-specific tasks. Option D (Scikit-learn) is for traditional machine learning, not transformer-based LLMs.

References:

NVIDIA NeMo Documentation: <https://docs.nvidia.com/deeplearning/nemo/user-guide/docs/en/stable/nlp/intro.html>

HuggingFace Transformers Documentation: <https://huggingface.co/docs/transformers/index>

**NEW QUESTION # 81**

What is the Open Neural Network Exchange (ONNX) format used for?

- A. Representing deep learning models
- B. Compressing deep learning models
- C. Reducing training time of neural networks
- D. Sharing neural network literature

**Answer: A**

Explanation:

The Open Neural Network Exchange (ONNX) format is an open-standard representation for deep learning models, enabling interoperability across different frameworks, as highlighted in NVIDIA's Generative AI and LLMs course. ONNX allows models trained in frameworks like PyTorch or TensorFlow to be exported and used in other compatible tools for inference or further development, ensuring portability and flexibility.

Option B is incorrect, as ONNX is not designed to reduce training time but to standardize model representation. Option C is wrong, as model compression is handled by techniques like quantization, not ONNX. Option D is inaccurate, as ONNX is unrelated to sharing literature. The course states: "ONNX is an open format for representing deep learning models, enabling seamless model exchange and deployment across various frameworks and platforms." References: NVIDIA Building Transformer-Based Natural Language Processing Applications course; NVIDIA Introduction to Transformer-Based Natural Language Processing.

**NEW QUESTION # 82**

What is the fundamental role of LangChain in an LLM workflow?

- A. To directly manage the hardware resources used by LLMs.
- B. To act as a replacement for traditional programming languages.
- C. To reduce the size of AI foundation models.
- D. To orchestrate LLM components into complex workflows.

**Answer: D**

Explanation:

LangChain is a framework designed to simplify the development of applications powered by large language models (LLMs) by orchestrating various components, such as LLMs, external data sources, memory, and tools, into cohesive workflows. According to NVIDIA's documentation on generative AI workflows, particularly in the context of integrating LLMs with external systems, LangChain enables developers to build complex applications by chaining together prompts, retrieval systems (e.g., for RAG), and memory modules to maintain context across interactions. For example, LangChain can integrate an LLM with a vector database for retrieval-augmented generation or manage conversational history for chatbots. Option A is incorrect, as LangChain complements, not replaces, programming languages. Option B is wrong, as LangChain does not modify model size. Option D is inaccurate, as hardware management is handled by platforms like NVIDIA Triton, not LangChain.

References:

NVIDIA NeMo Documentation: <https://docs.nvidia.com/deeplearning/nemo/user-guide/docs/en/stable/nlp/intro.html> LangChain Official Documentation: [https://python.langchain.com/docs/get\\_started/introduction](https://python.langchain.com/docs/get_started/introduction)

### NEW QUESTION # 83

Which of the following prompt engineering techniques is most effective for improving an LLM's performance on multi-step reasoning tasks?

- A. Chain-of-thought prompting with explicit intermediate steps.
- B. Zero-shot prompting with detailed task descriptions.
- C. Few-shot prompting with unrelated examples.
- D. Retrieval-augmented generation without context

Answer: A

Explanation:

Chain-of-thought (CoT) prompting is a highly effective technique for improving large language model (LLM) performance on multi-step reasoning tasks. By including explicit intermediate steps in the prompt, CoT guides the model to break down complex problems into manageable parts, improving reasoning accuracy. NVIDIA's NeMo documentation on prompt engineering highlights CoT as a powerful method for tasks like mathematical reasoning or logical problem-solving, as it leverages the model's ability to follow structured reasoning paths. Option A is incorrect, as retrieval-augmented generation (RAG) without context is less effective for reasoning tasks. Option B is wrong, as unrelated examples in few-shot prompting do not aid reasoning. Option C (zero-shot prompting) is less effective than CoT for complex reasoning.

References:

NVIDIA NeMo Documentation: <https://docs.nvidia.com/deeplearning/nemo/user-guide/docs/en/stable/nlp/intro.html> Wei, J., et al. (2022). "Chain-of-Thought Prompting Elicits Reasoning in Large Language Models."

### NEW QUESTION # 84

In the evaluation of Natural Language Processing (NLP) systems, what do 'validity' and 'reliability' imply regarding the selection of evaluation metrics?

- A. Validity involves the metric's ability to predict future trends in data, and reliability refers to its capacity to integrate with multiple data sources.
- B. Validity is concerned with the metric's computational cost, while reliability is about its applicability across different NLP platforms.
- C. Validity refers to the speed of metric computation, whereas reliability pertains to the metric's performance in high-volume data processing.
- D. Validity ensures the metric accurately reflects the intended property to measure, while reliability ensures consistent results over repeated measurements.

Answer: D

Explanation:

In evaluating NLP systems, as discussed in NVIDIA's Generative AI and LLMs course, validity and reliability are critical for selecting evaluation metrics. Validity ensures that a metric accurately measures the intended property (e.g., BLEU for translation quality or F1-score for classification performance), reflecting the system's true capability. Reliability ensures that the metric produces consistent results across repeated measurements under similar conditions, indicating stability and robustness. Together, these ensure trustworthy evaluations. Option A is incorrect, as validity is not about predicting trends, and reliability is not about data source integration. Option C is wrong, as validity and reliability are not primarily about computational cost or platform applicability. Option D is inaccurate, as validity and reliability do not focus on computation speed or high-volume processing. The course notes: "Validity ensures NLP evaluation metrics accurately measure the intended property, while reliability ensures consistent results across repeated evaluations, critical for robust system assessment." References: NVIDIA Building Transformer-Based Natural Language Processing Applications course; NVIDIA Introduction to Transformer-Based Natural Language Processing.

### NEW QUESTION # 85

.....

The goal of NCA-GENL preparation material is to help applicants prepare for the NVIDIA Generative AI LLMs certification exam by providing them with the Actual NCA-GENL Exam Questions they need to pass the exam. This NVIDIA Generative AI LLMs (NCA-GENL) study material is in the form of practice tests and NCA-GENL PDF that thoroughly covers the content of the test.

**NCA-GENL Dump Check:** [https://www.dumpkiller.com/NCA-GENL\\_braindumps.html](https://www.dumpkiller.com/NCA-GENL_braindumps.html)

- NCA-GENL Valid Dumps Files □ NCA-GENL Test Vce Free □ NCA-GENL Test Vce Free □ Copy URL ( [www.practicevce.com](http://www.practicevce.com) ) open and search for ( NCA-GENL ) to download for free □ NCA-GENL Valid Exam Topics
- NCA-GENL Latest Braindumps Ebook □ NCA-GENL Valid Test Tips □ NCA-GENL Free Updates □ Enter “ [www.pdfvce.com](http://www.pdfvce.com) ” and search for □ NCA-GENL □ to download for free □ Exam NCA-GENL Overviews
- Free NVIDIA NCA-GENL Exam Questions Updates and Demos □ Search for ▷ NCA-GENL ▲ and download exam materials for free through ✓ [www.troyecdumps.com](http://www.troyecdumps.com) □ ✓ □ □ Reliable NCA-GENL Exam Bootcamp
- Exam NCA-GENL Overviews □ Reliable NCA-GENL Exam Bootcamp □ NCA-GENL Test Pattern □ Open ➔ [www.pdfvce.com](http://www.pdfvce.com) □ □ □ enter ➔ NCA-GENL ▲ and obtain a free download □ NCA-GENL Test Vce Free
- Quiz NVIDIA - Newest NCA-GENL - NVIDIA Generative AI LLMs Exam Details □ Open website 《 [www.examdiscuss.com](http://www.examdiscuss.com) 》 and search for ( NCA-GENL ) for free download □ NCA-GENL Valid Test Tips
- High-quality NCA-GENL Exam Details bring you Correct NCA-GENL Dump Check for NVIDIA NVIDIA Generative AI LLMs □ 「 [www.pdfvce.com](http://www.pdfvce.com) 」 is best website to obtain 「 NCA-GENL 」 for free download □ NCA-GENL Latest Exam Review
- Free NVIDIA NCA-GENL Exam Questions Updates and Demos □ Easily obtain free download of ▷ NCA-GENL ▲ by searching on ✓ [www.exam4labs.com](http://www.exam4labs.com) □ ✓ □ □ NCA-GENL Test Vce Free
- Reliable NCA-GENL Exam Bootcamp □ NCA-GENL Sample Questions □ NCA-GENL Valid Exam Topics □ The page for free download of { NCA-GENL } on ➔ [www.pdfvce.com](http://www.pdfvce.com) □ will open immediately □ Valid NCA-GENL Test Duration
- Pass Guaranteed Quiz NVIDIA - Trustable NCA-GENL Exam Details □ Easily obtain ➔ NCA-GENL □ for free download through □ [www.validtorrent.com](http://www.validtorrent.com) □ □ NCA-GENL Sample Questions
- Exam NCA-GENL Overviews □ NCA-GENL Valid Exam Topics □ NCA-GENL Online Tests □ Search for ➔ NCA-GENL □ and download exam materials for free through ▷ [www.pdfvce.com](http://www.pdfvce.com) ▲ □ NCA-GENL Valid Dumps Files
- NCA-GENL Latest Braindumps Ebook □ NCA-GENL Valid Exam Topics □ Exam NCA-GENL Overviews □ Copy URL 【 [www.practicevce.com](http://www.practicevce.com) 】 open and search for 【 NCA-GENL 】 to download for free □ NCA-GENL Latest Test Fee
- [lms.ait.edu.za](http://lms.ait.edu.za), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [training.ifsinstitute.com](http://training.ifsinstitute.com), [hhi.instructure.com](http://hhi.instructure.com), [daotao.wisebusiness.edu.vn](http://daotao.wisebusiness.edu.vn), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [www.quora.com](http://www.quora.com), [user.xiaozhongwenhua.top](http://user.xiaozhongwenhua.top), [rdguitar.com](http://rdguitar.com), Disposable vapes

DOWNLOAD the newest Dumpkiller NCA-GENL PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1QM7-TC2LmR18BzeojMMIM4DgMhq0MKrC>