実用的NCA-GENL｜正確的なNCA-GENL合格対策試験｜試験の準備方法NVIDIA Generative AI LLMs問題集

無料でクラウドストレージから最新のMogiExam NCA-GENL PDFダンプをダウンロードする：https://drive.google.com/open?id=1U8DrhiwhQ_K-8MYn8d20kU8ikW81czAS

MogiExamはIT技術を勉強している人がよく知っているウェブサイトです。このサイトはIT認定試験を受けた受験生から広く好評されました。これはあなたに本当のヘルプを与えるサイトです。では、なぜMogiExamは皆さんの信頼を得ることができますか。それはMogiExamにはIT業界のエリートのグループがあって、グループのIT専門家達がずっと皆さんに最高のNCA-GENL資料を提供することに力を尽くしていますから。したがって、MogiExamは優れた参考書を提供して、みなさんのニーズを満たすことができます。

NCA-GENL試験ガイドには、高品質のサービスがあります。 24時間のオンラインサービスを提供しています。 NCA-GENL試験問題の使用中に質問がある場合は、メールでお問い合わせください。 私たちNVIDIAは最大限の忍耐と態度で優れたアフターサービスを提供します。 そして、NCA-GENL練習用トレントの使用中に発生する問題に対する詳細な解決策を提供します。 そして、NCA-GENLの学習資料は、あなたの監督と批判を歓迎します。 私たちのNCA-GENL学習教材の会社で、あなたはNVIDIA Generative AI LLMs成功の方向を見つけるでしょう。

>> NCA-GENL合格対策 <<

NCA-GENL問題集 & NCA-GENL日本語試験情報

なぜ受験生のほとんどはMogiExamを選んだのですか。それはMogiExamがすごく便利で、広い通用性があるからです。MogiExamのITエリートたちは彼らの専門的な目で、最新的なNVIDIAのNCA-GENL試験トレーニング資料に注目していて、うちのNVIDIAのNCA-GENL問題集の高い正確性を保証するのです。もし君はいささかな心配することがあるなら、あなたはうちの商品を購入する前に、MogiExamは無料でサンプルを提供することができます。

NVIDIA NCA-GENL 認定試験の出題範囲：

トピック	出題範囲
トピック 1	• 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.
トピック 2	• 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.

トピック 3	• 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.:
トピック 4	• 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.
トピック 5	• 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.

NVIDIA Generative AI LLMs 認定 NCA-GENL 試験問題 (Q63-Q68):

質問 # 63
Why is layer normalization important in transformer architectures?

- A. To stabilize the learning process by adjusting the inputs across the features.
- B. To compress the model size for efficient storage.
- C. To encode positional information within the sequence.
- D. To enhance the model's ability to generalize to new data.

正解：A

解説：
Layer normalization is a critical technique in Transformer architectures, as highlighted in NVIDIA's Generative AI and LLMs course. It stabilizes the learning process by normalizing the inputs to each layer across the features, ensuring that the mean and variance of the activations remain consistent. This is achieved by computing the mean and standard deviation of the inputs to a layer and scaling them to a standard range, which helps mitigate issues like vanishing or exploding gradients during training. This stabilization improves training efficiency and model performance, particularly in deep networks like Transformers. Option A is incorrect, as layer normalization primarily aids training stability, not generalization to new data, which is influenced by other factors like regularization. Option B is wrong, as layer normalization does not compress model size but adjusts activations. Option D is inaccurate, as positional information is handled by positional encoding, not layer normalization. The course notes: "Layer normalization stabilizes the training of Transformer models by normalizing layer inputs, ensuring consistent activation distributions and improving convergence." References: NVIDIA Building Transformer-Based Natural Language Processing Applications course; NVIDIA Introduction to Transformer-Based Natural Language Processing.

質問 # 64
In evaluating the transformer model for translation tasks, what is a common approach to assess its performance?

- A. Analyzing the lexical diversity of the model's translations compared to source texts.
- B. Comparing the model's output with human-generated translations on a standard dataset.
- C. Measuring the syntactic complexity of the model's translations against a corpus of professional translations.
- D. Evaluating the consistency of translation tone and style across different genres of text.

正解：B

解説：
A common approach to evaluate Transformer models for translation tasks, as highlighted in NVIDIA's Generative AI and LLMs course, is to compare the model's output with human-generated translations on a standard dataset, such as WMT (Workshop on Machine Translation) or BLEU-evaluated corpora. Metrics like BLEU (Bilingual Evaluation Understudy) score are used to quantify the similarity between machine and human translations, assessing accuracy and fluency. This method ensures objective, standardized evaluation.
Option A is incorrect, as lexical diversity is not a primary evaluation metric for translation quality. Option C is wrong, as tone and style consistency are secondary to accuracy and fluency. Option D is inaccurate, as syntactic complexity is not a standard evaluation criterion compared to direct human translation benchmarks.
The course states: "Evaluating Transformer models for translation involves comparing their outputs to human- generated translations on standard datasets, using metrics like BLEU to measure performance." References: NVIDIA Building Transformer-Based Natural

Language Processing Applications course; NVIDIA Introduction to Transformer-Based Natural Language Processing.

質問 # 65
Which calculation is most commonly used to measure the semantic closeness of two text passages?

- A. Jaccard similarity
- B. Hamming distance
- C. Euclidean distance
- D. Cosine similarity

正解：D

解說：
Cosine similarity is the most commonly used metric to measure the semantic closeness of two text passages in NLP. It calculates the cosine of the angle between two vectors (e.g., word embeddings or sentence embeddings) in a high-dimensional space, focusing on the direction rather than magnitude, which makes it robust for comparing semantic similarity. NVIDIA's documentation on NLP tasks, particularly in NeMo and embedding models, highlights cosine similarity as the standard metric for tasks like semantic search or text similarity, often using embeddings from models like BERT or Sentence-BERT. Option A (Hamming distance) is for binary data, not text embeddings. Option B (Jaccard similarity) is for set-based comparisons, not semantic content. Option D (Euclidean distance) is less common for text due to its sensitivity to vector magnitude.
References:
NVIDIA NeMo Documentation: https://docs.nvidia.com/deeplearning/nemo/user-guide/docs/en/stable/nlp/intro.html

質問 # 66
What metrics would you use to evaluate the performance of a RAG workflow in terms of the accuracy of responses generated in relation to the input query? (Choose two.)

- A. Retriever latency
- B. Generator latency
- C. Context precision
- D. Tokens generated per second
- E. Response relevancy

正解：C、E

解說：
In a Retrieval-Augmented Generation (RAG) workflow, evaluating the accuracy of responses relative to the input query focuses on the quality of the retrieved context and the generated output. As covered in NVIDIA's Generative AI and LLMs course, two key metrics are response relevancy and context precision. Response relevancy measures how well the generated response aligns with the input query, often assessed through human evaluation or automated metrics like ROUGE or BLEU, ensuring the output is pertinent and accurate.
Context precision evaluates the retriever's ability to fetch relevant documents or passages from the knowledge base, typically measured by metrics like precision@k, which assesses the proportion of retrieved items that are relevant to the query. Options A (generator latency), B (retriever latency), and C (tokens generated per second) are incorrect, as they measure performance efficiency (speed) rather than accuracy. The course notes:
"In RAG workflows, response relevancy ensures the generated output matches the query intent, while context precision evaluates the accuracy of retrieved documents, critical for high-quality responses." References: NVIDIA Building Transformer-Based Natural Language Processing Applications course; NVIDIA Introduction to Transformer-Based Natural Language Processing.

質問 # 67
In Exploratory Data Analysis (EDA) for Natural Language Understanding (NLU), which method is essential for understanding the contextual relationship between words in textual data?

- A. Computing the frequency of individual words to identify the most common terms in a text.
- B. Creating n-gram models to analyze patterns of word sequences like bigrams and trigrams.
- C. Applying sentiment analysis to gauge the overall sentiment expressed in a text.
- D. Generating word clouds to visually represent word frequency and highlight key terms.

正解：B

解説：
In Exploratory Data Analysis (EDA) for Natural Language Understanding (NLU), creating n-gram models is essential for understanding the contextual relationships between words, as highlighted in NVIDIA's Generative AI and LLMs course. N-grams (e.g., bigrams, trigrams) capture sequences of words, revealing patterns and dependencies in text, such as common phrases or syntactic structures, which are critical for NLU tasks like text generation or classification. Unlike single-word frequency analysis, n-grams provide insight into how words relate to each other in context. Option A is incorrect, as computing word frequencies focuses on individual terms, missing contextual relationships. Option B is wrong, as sentiment analysis targets overall text sentiment, not word relationships. Option C is inaccurate, as word clouds visualize frequency, not contextual patterns. The course notes: "N-gram models are used in EDA for NLU to analyze word sequence patterns, such as bigrams and trigrams, to understand contextual relationships in textual data." References: NVIDIA Building Transformer-Based Natural Language Processing Applications course; NVIDIA Introduction to Transformer-Based Natural Language Processing.

質問 # 68
......