# NCA-GENL模擬練習、NCA-GENL試験解説問題


BONUS！！！ Pass4Test NCA-GENLダンプの一部を無料でダウンロード：https://drive.google.com/open?id=1FbHt0fttNQFsM4fcP8QX-9qc6uU1jE3s

他の人はあちこちでNVIDIA　NCA-GENL試験資料を探しているとき、あなたはすでに勉強中で、準備階段でライバルに先立ちます。また、我々Pass4Testは量豊かのNVIDIA　NCA-GENL試験資料を提供しますし、ソフト版であなたにNVIDIA　NCA-GENL試験の最も現実的な環境をシミュレートさせます。勉強中で、何の質問があると、メールで我々はあなたのためにすぐ解決します。心配はありませんし、一心不乱に試験復習に取り組んでいます。

## 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 | • Prompt engineering: Focuses on techniques for designing and refining input prompts to effectively guide LLM outputs toward desired results. |
| トピック 2 | • Fundamentals of machine learning and neural networks: Covers the core concepts of how machine learning models learn from data, including the structure and function of neural networks that underpin large language models. |
| トピック 3 | • Data preprocessing and feature engineering: Covers preparing raw data through cleaning, transformation, and feature selection to make it suitable for model training. |

| | |
|---|---|
| トピック 4 | • Experimentation: Explores running and evaluating trials to test model behavior, compare approaches, and validate generative AI solutions. |
| トピック 5 | • Software development: Covers the programming practices and coding skills required to build, maintain, and deploy generative AI applications. |

**>> NCA-GENL模擬練習 <<**

# 完璧なNCA-GENL模擬練習 & 合格スムーズNCA-GENL試験解説問題 | 一生懸命にNCA-GENL関連合格問題

多くの時間とお金がいらなくて２０時間だけあって楽に一回にNVIDIAのNCA-GENL認定試験を合格できます。Pass4Testが提供したNVIDIAのNCA-GENL試験問題と解答が真実の試験の練習問題と解答は最高の相似性があります。

## NVIDIA Generative AI LLMs 認定 NCA-GENL 試験問題 (Q16-Q21):

**質問 # 16**
In the context of preparing a multilingual dataset for fine-tuning an LLM, which preprocessing technique is most effective for handling text from diverse scripts (e.g., Latin, Cyrillic, Devanagari) to ensure consistent model performance?

- A. Applying Unicode normalization to standardize character encodings.
- B. Removing all non-Latin characters to simplify the input.
- C. Converting text to phonetic representations for cross-lingual alignment.
- D. Normalizing all text to a single script using transliteration.

**正解：A**

解説：
When preparing a multilingual dataset for fine-tuning an LLM, applying Unicode normalization (e.g., NFKC or NFC forms) is the most effective preprocessing technique to handle text from diverse scripts like Latin, Cyrillic, or Devanagari. Unicode normalization standardizes character encodings, ensuring that visually identical characters (e.g., precomposed vs. decomposed forms) are represented consistently, which improves model performance across languages. NVIDIA's NeMo documentation on multilingual NLP preprocessing recommends Unicode normalization to address encoding inconsistencies in diverse datasets. Option A (transliteration) may lose linguistic nuances. Option C (removing non-Latin characters) discards critical information. Option D (phonetic conversion) is impractical for text-based LLMs.
References:
NVIDIA NeMo Documentation: https://docs.nvidia.com/deeplearning/nemo/user-guide/docs/en/stable/nlp/intro.html

**質問 # 17**
Which aspect in the development of ethical AI systems ensures they align with societal values and norms?

- A. Implementing complex algorithms to enhance AI's problem-solving capabilities.
- B. Ensuring AI systems have explicable decision-making processes.
- C. Developing AI systems with autonomy from human decision-making.
- D. Achieving the highest possible level of prediction accuracy in AI models.

**正解：B**

解説：
Ensuring explicable decision-making processes, often referred to as explainability or interpretability, is critical for aligning AI systems with societal values and norms. NVIDIA's Trustworthy AI framework emphasizes that explainable AI allows stakeholders to understand how decisions are made, fostering trust and ensuring compliance with ethical standards. This is particularly important for addressing biases and ensuring fairness. Option A (prediction accuracy) is important but does not guarantee ethical alignment. Option B (complex algorithms) may improve performance but not societal alignment. Option C (autonomy) can conflict with ethical

oversight, making it less desirable.
References:
NVIDIA Trustworthy AI:https://www.nvidia.com/en-us/ai-data-science/trustworthy-ai/

**質問 # 18**
Which of the following prompt engineering techniques is most effective for improving an LLM's performance on multi-step reasoning tasks?

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

**正解：C**

解説：
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."

**質問 # 19**
In the Transformer architecture, which of the following statements about the Q (query), K (key), and V (value) matrices is correct?

- A. K is responsible for computing the attention scores between the query and key vectors.
- B. V is used to calculate the positional embeddings for each token in the input sequence.
- C. Q represents the query vector used to retrieve relevant information from the input sequence.
- D. Q, K, and V are randomly initialized weight matrices used for positional encoding.

**正解：C**

解説：
In the transformer architecture, the Q (query), K (key), and V (value) matrices are used in the self-attention mechanism to compute relationships between tokens in a sequence. According to "Attention is All You Need" (Vaswani et al., 2017) and NVIDIA's NeMo documentation, the query vector (Q) represents the token seeking relevant information, the key vector (K) is used to compute compatibility with other tokens, and the value vector (V) provides the information to be retrieved. The attention score is calculated as a scaled dot- product of Q and K, and the output is a weighted sum of V. Option C is correct, as Q retrieves relevant information. Option A is incorrect, as Q, K, and V are not used for positional encoding. Option B is wrong, as attention scores are computed using both Q and K, not K alone. Option D is false, as positional embeddings are separate from V.
References:
Vaswani, A., et al. (2017). "Attention is All You Need."
NVIDIA NeMo Documentation:https://docs.nvidia.com/deeplearning/nemo/user-guide/docs/en/stable/nlp
/intro.html

**質問 # 20**
Which of the following best describes Word2vec?

- A. A deep learning algorithm used to generate word embeddings from text data.
- B. A statistical technique used to analyze word frequency in a text corpus.
- C. A database management system designed for storing and querying word data.

- D. A programming language used to build artificial intelligence models.

**正解：A**

解説：
Word2Vec is a groundbreaking deep learning algorithm developed to create dense vector representations, or embeddings, of words based on their contextual usage in large text corpora. Unlike traditional methods like bag-of-words or TF-IDF, which rely on frequency counts and often result in sparse vectors, Word2Vec employs neural networks to learn continuous vector spaces where semantically similar words are positioned closer together. This enables machines to capture nuances such as synonyms, analogies, and relationships (e.
g., "king" - "man" + "woman" # "queen"). The algorithm operates through two primary architectures:
Continuous Bag-of-Words (CBOW), which predicts a target word from its surrounding context, and Skip- Gram, which does the reverse by predicting context words from a target word. Skip-Gram is particularly effective for rare words and larger datasets, while CBOW is faster and better for frequent words. In the context of NVIDIA's Generative AI and LLMs course, Word2Vec is highlighted as a foundational step in the evolution of text embeddings in natural language processing (NLP) tasks, paving the way for more advanced models like RNN-based embeddings and Transformers. This is essential for understanding how LLMs build upon these embeddings for tasks such as semantic analysis and language generation. Exact extract from the course description:
"Understand how text embeddings have rapidly evolved in NLP tasks such as Word2Vec, recurrent neural network (RNN)-based embeddings, and Transformers." This positions Word2Vec as a key deep learning technique for generating meaningful word vectors from text data, distinguishing it from mere statistical frequency analysis or unrelated tools like programming languages or databases

**質問 # 21**
.....

- mekkawyacademy.com, www.evstudy.com, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, bbs.t-firefly.com, startupxplore.com, bbs.t-firefly.com, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, Disposable vapes

P.S. Pass4TestがGoogle Driveで共有している無料かつ新しいNCA-GENLダンプ：https://drive.google.com/open?id=1FbHt0fttNQFsM4fcP8QX-9qc6uU1jE3s