

NCA-GENL套裝，NCA-GENL考古題推薦



從Google Drive中免費下載最新的KaoGuTi NCA-GENL PDF版考試題庫：<https://drive.google.com/open?id=17MM0PyYoSYtrP0knskeKHKtaDQCXDuwZ>

通過那些很多已經通過NVIDIA NCA-GENL 認證考試的IT專業人員的回饋，他們的成功得益於KaoGuTi的說明。KaoGuTi提供的針對性測試練習題和答案給了他們很大幫助，節約了他們的寶貴的時間和精力，讓他們輕鬆順利地通過他們第一次參加的NVIDIA NCA-GENL 認證考試。所以KaoGuTi是個值得你們信賴的網站。選擇了KaoGuTi，下一個成功的IT人士就是你，KaoGuTi會成就你的夢想。

NVIDIA NCA-GENL 考試大綱：

主題	簡介
主題 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.
主題 2	<ul style="list-style-type: none">Experiment Design
主題 3	<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.
主題 4	<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.
主題 5	<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.
主題 6	<ul style="list-style-type: none">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.

主題 7	<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.
主題 8	<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.
主題 9	<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.
主題 10	<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.:

>> NCA-GENL套裝 <<

免費下載NCA-GENL套裝和資格考試中的領先供應商和值得信賴的NCA-GENL考古題推薦

為通過NVIDIA NCA-GENL認證考試花大量的時間和精力復習相關知識，但是卻是冒險地通過考試。選擇KaoGuTi的產品卻可以讓你花少量的錢，一次性安全通過考試。我相信在如今時間如此寶貴的社會裏，KaoGuTi更適合你的選擇。而且我們的KaoGuTi是眾多類似網站中最能給你保障的一個網站，選擇KaoGuTi就等於選擇了成功。

最新的 NVIDIA-Certified Associate NCA-GENL 免費考試真題 (Q36-Q41):

問題 #36

What statement best describes the diffusion models in generative AI?

- A. Diffusion models are generative models that use a transformer architecture to learn the underlying probability distribution of the data.
- B. Diffusion models are probabilistic generative models that progressively inject noise into data, then learn to reverse this process for sample generation.
- C. Diffusion models are unsupervised models that use clustering algorithms to group similar data points together.
- D. Diffusion models are discriminative models that use gradient-based optimization algorithms to classify data points.

答案: B

解題說明:

Diffusion models, as discussed in NVIDIA's Generative AI and LLMs course, are probabilistic generative models that operate by progressively adding noise to data in a forward process and then learning to reverse this process to generate new samples. This involves a Markov chain that gradually corrupts data with noise and a reverse process that denoises it to reconstruct realistic samples, making them powerful for generating high-quality images, text, and other data. Unlike Transformer-based models, diffusion models rely on this iterative denoising mechanism. Option B is incorrect, as diffusion models are generative, not discriminative, and focus on data generation, not classification. Option C is wrong, as diffusion models do not use clustering algorithms but focus on generative tasks. Option D is inaccurate, as diffusion models do not inherently rely on Transformer architectures but use distinct denoising processes. The course states: "Diffusion models are probabilistic generative models that add noise to data and learn to reverse the process for sample generation, widely used in generative AI tasks." References: NVIDIA Building Transformer-Based Natural Language Processing Applications course; NVIDIA Introduction to Transformer-Based Natural Language Processing.

問題 #37

What is a foundation model in the context of Large Language Models (LLMs)?

- A. Any model based on the foundation paper "Attention is all you need," that uses recurrent neural networks and convolution

layers.

- B. Any model trained on vast quantities of data at scale whose goal is to serve as a starter that can be adapted to a variety of downstream tasks.
- C. A model that sets the state-of-the-art results for any of the tasks that compose the General Language Understanding Evaluation (GLUE) benchmark.
- D. Any model validated by the artificial intelligence safety institute as the foundation for building transformer-based applications.

答案： B

解題說明：

In the context of Large Language Models (LLMs), a foundation model refers to a large-scale model trained on vast quantities of diverse data, designed to serve as a versatile starting point that can be fine-tuned or adapted for a variety of downstream tasks, such as text generation, classification, or translation. As covered in NVIDIA's Generative AI and LLMs course, foundation models like BERT, GPT, or T5 are pre-trained on massive datasets and can be customized for specific applications, making them highly flexible and efficient.

Option A is incorrect, as achieving state-of-the-art results on GLUE is not a defining characteristic of foundation models, though some may perform well on such benchmarks. Option C is wrong, as there is no specific validation by an AI safety institute required to define a foundation model. Option D is inaccurate, as the "Attention is All You Need" paper introduced Transformers, which rely on attention mechanisms, not recurrent neural networks or convolution layers. The course states: "Foundation models are large-scale models trained on broad datasets, serving as a base for adaptation to various downstream tasks in NLP." References: NVIDIA Building Transformer-Based Natural Language Processing Applications course; NVIDIA Introduction to Transformer-Based Natural Language Processing.

問題 #38

What is the prompt "Translate English to French: cheese =>" an example of?

- A. Few-shot learning
- B. Fine tuning a model
- C. One-shot learning
- D. Zero-shot learning

答案： D

解題說明：

The prompt "Translate English to French: cheese =>" is an example of zero-shot learning, as discussed in NVIDIA's Generative AI and LLMs course. Zero-shot learning refers to a model's ability to perform a task without prior task-specific training or examples, relying solely on its pre-trained knowledge and the prompt's instructions. In this case, the prompt provides no training examples, expecting the model to translate "cheese" to French ("fromage") based on its general understanding of language and translation. Option A, few-shot learning, is incorrect, as it involves providing a few examples in the prompt. Option B, fine-tuning, involves retraining the model, not prompting. Option C, one-shot learning, requires a single example, which is not provided here. The course notes: "Zero-shot learning enables LLMs to perform tasks like translation without task-specific training, using only a descriptive prompt to leverage pre-trained knowledge." References: NVIDIA Building Transformer-Based Natural Language Processing Applications course; NVIDIA Introduction to Transformer-Based Natural Language Processing.

問題 #39

When designing prompts for a large language model to perform a complex reasoning task, such as solving a multi-step mathematical problem, which advanced prompt engineering technique is most effective in ensuring robust performance across diverse inputs?

- A. Retrieval-augmented generation with external mathematical databases.
- B. Chain-of-thought prompting with step-by-step reasoning examples.
- C. Zero-shot prompting with a generic task description.
- D. Few-shot prompting with randomly selected examples.

答案： B

解題說明：

Chain-of-thought (CoT) prompting is an advanced prompt engineering technique that significantly enhances a large language model's (LLM) performance on complex reasoning tasks, such as multi-step mathematical problems. By including examples that explicitly

demonstrate step-by-step reasoning in the prompt, CoT guides the model to break down the problem into intermediate steps, improving accuracy and robustness.

NVIDIA's NeMo documentation on prompt engineering highlights CoT as a powerful method for tasks requiring logical or sequential reasoning, as it leverages the model's ability to mimic structured problem-solving. Research by Wei et al. (2022) demonstrates that CoT outperforms other methods for mathematical reasoning. Option A (zero-shot) is less effective for complex tasks due to lack of guidance. Option B (few-shot with random examples) is suboptimal without structured reasoning. Option D (RAG) is useful for factual queries but less relevant for pure reasoning tasks.

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."

問題 #40

Which of the following is an activation function used in neural networks?

- A. Diffusion function
- B. **Sigmoid function**
- C. K-means clustering function
- D. Mean Squared Error function

答案: B

解題說明:

The sigmoid function is a widely used activation function in neural networks, as covered in NVIDIA's Generative AI and LLMs course. It maps input values to a range between 0 and 1, making it particularly useful for binary classification tasks and as a non-linear activation in early neural network architectures. The sigmoid function, defined as $f(x) = 1 / (1 + e^{-x})$

P.S. KaoGuTi在Google Drive上分享了免費的、最新的NCA-GENL考試題庫: <https://drive.google.com/open?id=17MM0PyYoSYtrP0knskeKHKitDQCXDuwZ>