Number: H13-311_V3.5
Passing Score: 800
Time Limit: 120
File Version: 3.0

Exam Code: H13-311_V3.5

Exam Name: HCIA-AI V3.5

Exam A

QUESTION 1
In MindSpore, mindspore.nn.Conv2d() is used to create a convolutional layer. Which of the following values can be passed to this API's 'pad_mode' parameter?

A. pad
B. same
C. valid
D. nopadding

Correct Answer: B, C
Section:
Explanation:
The pad_mode parameter in mindspore.nn.Conv2d() can take values such as:
same: Ensures the output has the same spatial dimensions as the input.
valid: Performs convolution without padding, resulting in an output smaller than the input.
Other values like 'pad' and 'nopadding' are not valid options for the pad_mode parameter.

QUESTION 2
As we understand more about machine learning, we will find that its scope is constantly changing over time.

A. TRUE
B. FALSE

Correct Answer: A
Section:
Explanation:
Machine learning is a rapidly evolving field, and its scope indeed changes over time. With advancements in computational power, the introduction of new algorithms, frameworks, and techniques, and the growing availability of data, the capabilities of machine learning have expanded significantly. Initially, machine learning was limited to simpler algorithms like linear regression, decision trees, and k-nearest neighbors. Over time, however, more complex approaches such as deep learning and reinforcement learning have emerged, dramatically increasing the applications and effectiveness of machine learning solutions.
In the Huawei HCIA-AI curriculum, it is emphasized that AI, especially machine learning, has become more powerful due to these continuous developments, allowing it to be applied to broader and more complex problems. The framework and methodologies in machine learning have evolved, making it possible to perform more sophisticated tasks such as real-time decision-making, image recognition, natural language processing, and even autonomous driving.
As technology advances, the scope of machine learning will continue to shift, providing new opportunities for innovation. This is why it is important to stay updated on recent developments to fully leverage machine learning in various AI applications.

QUESTION 3
Which of the following is the order of tensor [[0,1],[2,3]]?

A. 6
B. 3
C. 2
D. 4

Correct Answer: C

Section:
Explanation:
The order of a tensor refers to its rank, which is the number of dimensions it has. For the tensor [[0,1],[2,3]], the rank is 2 because it is a 2x2 matrix, meaning it has 2 dimensions.

QUESTION 4
When you use MindSpore to execute the following code, which of the following is the output?
```
from mindspore import ops
import mindspore
shape = (2, 2)
ones = ops.Ones()
output = ones(shape, dtype=mindspore.float32)
print(output)
```

A. [[1 1] [1 1]]
B. [[1. 1.] [1. 1.]]
C. 1
D. [[1. 1. 1. 1.]]

Correct Answer: B
Section:
Explanation:
In MindSpore, using ops.Ones() with a specified shape and dtype=mindspore.float32 will create a tensor of ones with floating-point values. The output will be a 2x2 matrix filled with 1.0 values. The floating-point format (with a decimal point) ensures that the output is in the form of [[1. 1.], [1. 1.]].

QUESTION 5
Huawei's full-stack AI solution includes Ascend, MindSpore, and ModelArts. (Enter an acronym.)

A. AII
B. AIIS
C. CANN
D. None of the above

Correct Answer: C
Section:
Explanation:
CANN (Compute Architecture for Neural Networks) is part of Huawei's full-stack AI solution, which includes Ascend (hardware), MindSpore (AI framework), and ModelArts (AI development platform). CANN optimizes the computing efficiency of AI models and provides basic software components for the Ascend AI processors. This architecture supports deep learning and machine learning tasks by enhancing computational performance and providing better neural network training efficiency.
Together, Ascend, MindSpore, and CANN form a critical infrastructure that underpins Huawei's AI development ecosystem, allowing seamless integration from hardware to software.

QUESTION 6
The concept of 'artificial intelligence' was first proposed in the year of:

A. 1950
B. 1956
C. 1960
D. 1965

Correct Answer: B
Section:
Explanation:
The concept of 'artificial intelligence' was first formally introduced in 1956 during the Dartmouth Conference, organized by John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon. This event is widely regarded as the birth of AI as a field of study. The conference aimed to explore the idea that human intelligence could be simulated by machines, laying the groundwork for subsequent AI research and development.
This date is significant in the history of AI because it marked the beginning of a concentrated effort to develop machines that could mimic cognitive functions such as learning, reasoning, and problem-solving.

QUESTION 7
Which of the following are subfields of AI?

A. Backpropagation algorithm
B. Expert system
C. Smart finance
D. Computer vision

Correct Answer: B, D
Section:
Explanation:
Artificial intelligence is a broad field that encompasses several subfields. Two key subfields are:
Expert systems, which are computer programs that mimic the decision-making abilities of a human expert by reasoning through bodies of knowledge. These systems are used in various domains such as healthcare, engineering, and finance.
Computer vision, which enables machines to interpret and understand visual data from the world. It includes tasks such as object detection, image recognition, and video analysis.
While options like backpropagation and smart finance are related to AI, they represent specific algorithms or application areas rather than broad subfields.

QUESTION 8
What are the application scenarios of computer vision?

A. Video action analysis
B. Image search
C. Smart albums
D. Voice navigation

Correct Answer: A, B, C
Section:
Explanation:
Computer vision, a subfield of AI, has various application scenarios that involve the analysis and understanding of images and videos. Some key application scenarios include:
Video action analysis: Identifying and analyzing human actions or movements in videos.
Image search: Using visual information to search for similar images in large databases.
Smart albums: Organizing and categorizing photos using AI-based image recognition algorithms to group them by themes, people, or events.
Voice navigation is a part of natural language processing and speech recognition, not computer vision.

QUESTION 9
Which of the following is NOT a commonly used AI computing framework?

A. PyTorch
B. MindSpore
C. TensorFlow
D. OpenCV

Correct Answer: D
Section:
Explanation:
OpenCV is a library used primarily for computer vision tasks like image and video processing. It is not considered an AI computing framework in the same way as PyTorch, MindSpore, or TensorFlow, which are commonly used frameworks for developing AI and machine learning models. AI frameworks like PyTorch, TensorFlow, and Huawei's MindSpore are designed to facilitate the development and deployment of deep learning models.

QUESTION 10
'Today's speech processing technology can achieve a recognition accuracy of over 90% in any case.' Which of the following is true about this statement?

A. This statement is incorrect. The accuracy of speech recognition is high, but not extremely high.
B. This statement is incorrect. In many situations, noise and background sound have a huge impact on speech recognition accuracy.
C. This statement is correct. Speech processing can achieve a high level of accuracy.
D. This statement is correct. Speech processing has a long history and the technology is very mature.

Correct Answer: B
Section:
Explanation:
While speech recognition technology has improved significantly, its accuracy can still be affected by external factors such as noise, background sound, accents, and speech clarity. Although systems can achieve over 90% accuracy under controlled conditions, the accuracy drops in noisy or complex real-world environments. Therefore, the statement that today's speech processing technology can always achieve high recognition accuracy is incorrect.
Speech recognition systems are sophisticated but still face challenges in environments with heavy noise, where the technology has difficulty interpreting speech accurately.

QUESTION 11
'AI application fields include only computer vision and speech processing.' Which of the following is true about this statement?

A. This statement is false. The application fields of AI include computer vision, speech processing, natural language processing, and others.
B. This statement is false. AI application fields include only computer vision and natural language processing.
C. This statement is true. Voice data is processed with extremely high accuracy.
D. This statement is true. Computer vision is the most important AI application.

Correct Answer: A
Section:
Explanation:
AI is not limited to just computer vision and speech processing. In addition to these fields, AI encompasses other important areas such as natural language processing (NLP), robotics, smart finance, autonomous driving, and more. Natural language processing focuses on understanding and generating human language, while other fields apply AI to various industries and applications such as healthcare, finance, and manufacturing. AI is a broad field with numerous application areas.

QUESTION 12
Which of the following are common gradient descent methods?

A. Batch gradient descent (BGD)
B. Mini-batch gradient descent (MBGD)
C. Multi-dimensional gradient descent (MDGD)
D. Stochastic gradient descent (SGD)

Correct Answer: A, B, D
Section:
Explanation:

The gradient descent method is a core optimization technique in machine learning, particularly for neural networks and deep learning models. The common gradient descent methods include:
Batch Gradient Descent (BGD): Updates the model parameters after computing the gradients from the entire dataset.
Mini-batch Gradient Descent (MBGD): Updates the model parameters using a small batch of data, combining the benefits of both batch and stochastic gradient descent.
Stochastic Gradient Descent (SGD): Updates the model parameters for each individual data point, leading to faster but noisier updates.
Multi-dimensional gradient descent is not a recognized method in AI or machine learning.

QUESTION 13
Which of the following algorithms presents the most chaotic landscape on the loss surface?

A. Stochastic gradient descent
B. MGD
C. MBGD
D. BGD

Correct Answer: A
Section:
Explanation:
Stochastic Gradient Descent (SGD) presents the most chaotic landscape on the loss surface because it updates the model parameters for each individual training example, which can introduce a significant amount of noise into the optimization process. This leads to a less smooth and more chaotic path toward the global minimum compared to methods like batch gradient descent or mini-batch gradient descent, which provide more stable updates.

QUESTION 14
Which of the following statements are true about the k-nearest neighbors (k-NN) algorithm?

A. k-NN typically uses the mean value method to predict regression.
B. k-NN typically uses the majority voting method to predict classification.
C. k-NN is a parametric method often used for datasets with regular decision boundaries.
D. The k-NN algorithm determines which class an object belongs to based on the class to which most of the object's k nearest neighbors belong.

Correct Answer: B, D
Section:
Explanation:
The k-nearest neighbors (k-NN) algorithm is a non-parametric algorithm used for both classification and regression. In classification tasks, it typically uses majority voting to assign a label to a new instance based on the most common class among its nearest neighbors. The algorithm works by calculating the distance (often using Euclidean distance) between the query point and the points in the dataset, and then assigning the query point to the class that is most frequent among its k nearest neighbors.
For regression tasks, k-NN can predict the outcome based on the mean of the values of the k nearest neighbors, although this is less common than its classification use.

QUESTION 15
An algorithm of unsupervised learning classifies samples in a dataset into several categories. Samples belonging to the same category have high similarity.

A. TRUE
B. FALSE

Correct Answer: A
Section:
Explanation:
In unsupervised learning, the goal is to find hidden patterns or intrinsic structures in input data without labeled outcomes. One common unsupervised learning task is clustering, where an algorithm groups the dataset into several categories or clusters. Samples within the same cluster have high similarity based on certain features, while samples in different clusters have low similarity. Examples of clustering algorithms include k-means and

hierarchical clustering.

QUESTION 16
Which of the following statements is false about the debugging and application of a regression model?

A. If the model does not meet expectations, you need to use data cleansing and feature engineering.
B. After model training is complete, you need to use the test dataset to evaluate your model so that its generalization capability meets expectations.
C. If overfitting occurs, you can add a regularization term to the Lasso or ridge regression and adjust hyperparameters.
D. If underfitting occurs, you can use a more complex regression model, for example, logistic regression.

Correct Answer: D
Section:
Explanation:
Logistic regression is not a solution for underfitting in regression models, as it is used primarily for classification problems rather than regression tasks. If underfitting occurs, it means that the model is too simple to capture the underlying patterns in the data. Solutions include using a more complex regression model like polynomial regression or increasing the number of features in the dataset.
Other options like adding a regularization term for overfitting (Lasso or Ridge) and using data cleansing and feature engineering are correct methods for improving model performance.

QUESTION 17
Which of the following statements is false about feedforward neural networks?

A. A unidirectional multi-layer structure is adopted. Each layer includes several neurons, and those in the same layer are not connected to each other. Only unidirectional inter-layer information transmission is supported.
B. Nodes at each hidden layer represent neurons that provide the computing function.
C. Input nodes do not provide the computing function and are used to represent only the element values of an input vector.
D. Each neuron is connected to all neurons at the previous layer.

Correct Answer: D
Section:
Explanation:
This statement is false because not all feedforward neural networks follow this architecture. While fully-connected layers do have this type of connectivity (where each neuron is connected to all neurons in the previous layer), feedforward networks can include layers like convolutional layers, where not every neuron is connected to all previous neurons. Convolutional layers, common in convolutional neural networks (CNNs), only connect to a local region of the input, preserving spatial information.

QUESTION 18
Which of the following are feedforward neural networks?

A. Fully-connected neural networks
B. Recurrent neural networks
C. Boltzmann machines
D. Convolutional neural networks

Correct Answer: A, D
Section:
Explanation:
Feedforward neural networks (FNNs) are networks where information moves in only one direction---forward---from the input nodes through hidden layers to the output nodes. Both fully-connected neural networks (where each neuron in one layer connects to every neuron in the next) and convolutional neural networks (CNNs) (which have a specific architecture for image data) are examples of feedforward networks.
However, recurrent neural networks (RNNs) and Boltzmann machines are not feedforward networks. RNNs include loops where information can be fed back into previous layers, and Boltzmann machines involve undirected connections between units, making them a form of a stochastic network rather than a feedforward structure.

QUESTION 19
The mean squared error (MSE) loss function cannot be used for classification problems.

A. TRUE
B. FALSE

Correct Answer: A
Section:
Explanation:
The mean squared error (MSE) loss function is primarily used for regression problems, where the goal is to minimize the difference between the predicted and actual continuous values. For classification problems, where the target output is categorical (e.g., binary or multi-class labels), loss functions like cross-entropy are more suitable, as they are designed to handle the probabilistic interpretation of outputs in classification tasks.
Using MSE for classification could lead to inefficient training because it doesn't capture the probabilistic relationships required for classification tasks.

QUESTION 20
Which of the following statements is false about gradient descent algorithms?

A. Each time the global gradient updates its weight, all training samples need to be calculated.
B. When GPUs are used for parallel computing, the mini-batch gradient descent (MBGD) takes less time than the stochastic gradient descent (SGD) to complete an epoch.
C. The global gradient descent is relatively stable, which helps the model converge to the global extremum.
D. When there are too many samples and GPUs are not used for parallel computing, the convergence process of the global gradient algorithm is time-consuming.

Correct Answer: B
Section:
Explanation:
The statement that mini-batch gradient descent (MBGD) takes less time than stochastic gradient descent (SGD) to complete an epoch when GPUs are used for parallel computing is incorrect. Here's why:
Stochastic Gradient Descent (SGD) updates the weights after each training sample, which can lead to faster updates but more noise in the gradient steps. It completes an epoch after processing all samples one by one.
Mini-batch Gradient Descent (MBGD) processes small batches of data at a time, updating the weights after each batch. While MBGD leverages the computational power of GPUs effectively for parallelization, the comparison made in this question is not about overall computation speed, but about completing an epoch.
MBGD does not necessarily complete an epoch faster than SGD, as MBGD processes multiple samples in each batch, meaning fewer updates per epoch compared to SGD, where weights are updated after every individual sample.
Therefore, the correct answer is B. FALSE, as MBGD does not always take less time than SGD for completing an epoch, even when GPUs are used for parallelization.
HCIA AI
AI Development Framework: Discussion of gradient descent algorithms and their efficiency on different hardware architectures like GPUs.

QUESTION 21
All kernels of the same convolutional layer in a convolutional neural network share a weight.

A. TRUE
B. FALSE

Correct Answer: B
Section:
Explanation:
In a convolutional neural network (CNN), each kernel (also called a filter) in the same convolutional layer does not share weights with other kernels. Each kernel is independent and learns different weights during training to detect different features in the input data. For instance, one kernel might learn to detect edges, while another might detect textures.
However, the same kernel's weights are shared across all spatial positions it moves across the input feature map. This concept of weight sharing is what makes CNNs efficient and well-suited for tasks like image recognition.
Thus, the statement that all kernels share weights is false.
HCIA AI
Deep Learning Overview: Detailed description of CNNs, focusing on kernel operations and weight sharing mechanisms within a single kernel, but not across different kernels.

QUESTION 22
The core of the MindSpore training data processing engine is to efficiently and flexibly convert training samples (datasets) to MindRecord and provide them to the training network for training.

A. TRUE
B. FALSE

Correct Answer: A
Section:
Explanation:
MindSpore, Huawei's AI framework, includes a data processing engine designed to efficiently handle large datasets during model training. The core feature of this engine is the ability to convert training samples into a format called MindRecord, which optimizes data input and output processes for training. This format ensures that the data pipeline is fast and flexible, providing data efficiently to the training network.
The statement is true because one of MindSpore's core functionalities is to preprocess data and optimize its flow into the neural network training pipeline using the MindRecord format.
HCIA AI
Introduction to Huawei AI Platforms: Covers MindSpore's architecture, including its data processing engine and the use of the MindRecord format for efficient data management.

QUESTION 23
When using the following code to construct a neural network, MindSpore can inherit the Cell class and rewrite the __init__ and construct methods.

A. TRUE
B. FALSE

Correct Answer: A
Section:
Explanation:
In MindSpore, the neural network structure is defined by inheriting the Cell class, which represents a computational node or a layer in the network. Users can customize the network by overriding the __init__ method (for initializing layers) and the construct method (for defining the forward pass of the network). This modular design allows for easy and flexible neural network construction.
Thus, the statement is true because MindSpore's framework allows developers to build neural networks by extending the Cell class and defining custom behavior through the __init__ and construct methods.
HCIA AI
AI Development Framework: Detailed coverage of building neural networks in MindSpore, including how to inherit from the Cell class and rewrite key methods for custom network architecture.

QUESTION 24
Which of the following is NOT a key feature that enables all-scenario deployment and collaboration for MindSpore?

A. Data and computing graphs are transmitted to Ascend AI Processors.
B. Federal meta-learning enables real-time, coordinated model updates between different devices, and across the device and cloud.
C. Unified model IR delivers a consistent deployment experience.
D. Graph optimization based on a software-hardware synergy shields the differences between scenarios.

Correct Answer: B
Section:
Explanation:
While MindSpore supports all-scenario deployment with features like data and computing graph transmission to Ascend AI processors, unified model IR for consistent deployment, and graph optimization based on software-hardware synergy, federal meta-learning is not explicitly a core feature of MindSpore's deployment strategy. Federal meta-learning refers to a distributed learning paradigm, but MindSpore focuses more on efficient computing and model optimization across different environments.

QUESTION 25
In MindSpore, the basic unit of the neural network is nn.Cell.

A. TRUE
B. FALSE

Correct Answer: A
Section:
Explanation:
In MindSpore, nn.Cell is the basic unit of a neural network. It represents layers, models, and other neural network components, encapsulating the forward logic of the network. It allows users to define, organize, and manage neural network layers in MindSpore, making it a core building block in neural network construction.

QUESTION 26
As the cornerstone of Huawei's full-stack, all-scenario AI solution, it provides modules, boards, and servers powered by the Ascend AI processor to meet customer demand for computing power in all scenarios.

A. Atlas
B. CANN
C. MindSpore
D. ModelArts

Correct Answer: A
Section:
Explanation:
Atlas is a key part of Huawei's full-stack, all-scenario AI solution. It provides AI hardware resources in the form of modules, boards, edge stations, and servers powered by Huawei's Ascend AI processors. The Atlas series is designed to meet customer demands for AI computing power in a variety of deployment scenarios, including cloud, edge, and device environments.
Huawei's full-stack AI solution aims to deliver comprehensive AI capabilities across different levels. The Atlas series supports a wide range of industries by offering scalable AI computing resources, which are critical for industries dealing with large volumes of data and needing high-performance computing.

QUESTION 27
Which of the following are callback options provided by MindSpore?

A. SummaryCollector
B. TrainStep
C. ModelCheckpoint
D. LossMonitor

Correct Answer: A, C, D
Section:
Explanation:
MindSpore provides several callback functions that can be used to monitor, modify, or control the behavior of the training process. These include:
SummaryCollector: Collects summaries such as loss and accuracy for visualization and monitoring.
ModelCheckpoint: Saves model parameters during or after training.
LossMonitor: Monitors the loss values during training and can stop training if certain conditions are met.
TrainStep is not a callback but rather a fundamental step in training.

QUESTION 28
Which of the following statements are false about softmax and logistic?

A. In terms of probability, softmax modeling uses the polynomial distribution, whereas logistic modeling uses the binomial distribution.
B. Multiple logistic regressions can be combined to achieve multi-class classification effects.
C. Logistic is used for binary classification problems, whereas softmax is used for multi-class classification problems.

D. In the multi-class classification of softmax regression, the output classes are not mutually exclusive. That is, the word 'Apple' belongs to both the 'fruit' and '3C' classes.

Correct Answer: A, D
Section: