Lead2passExam

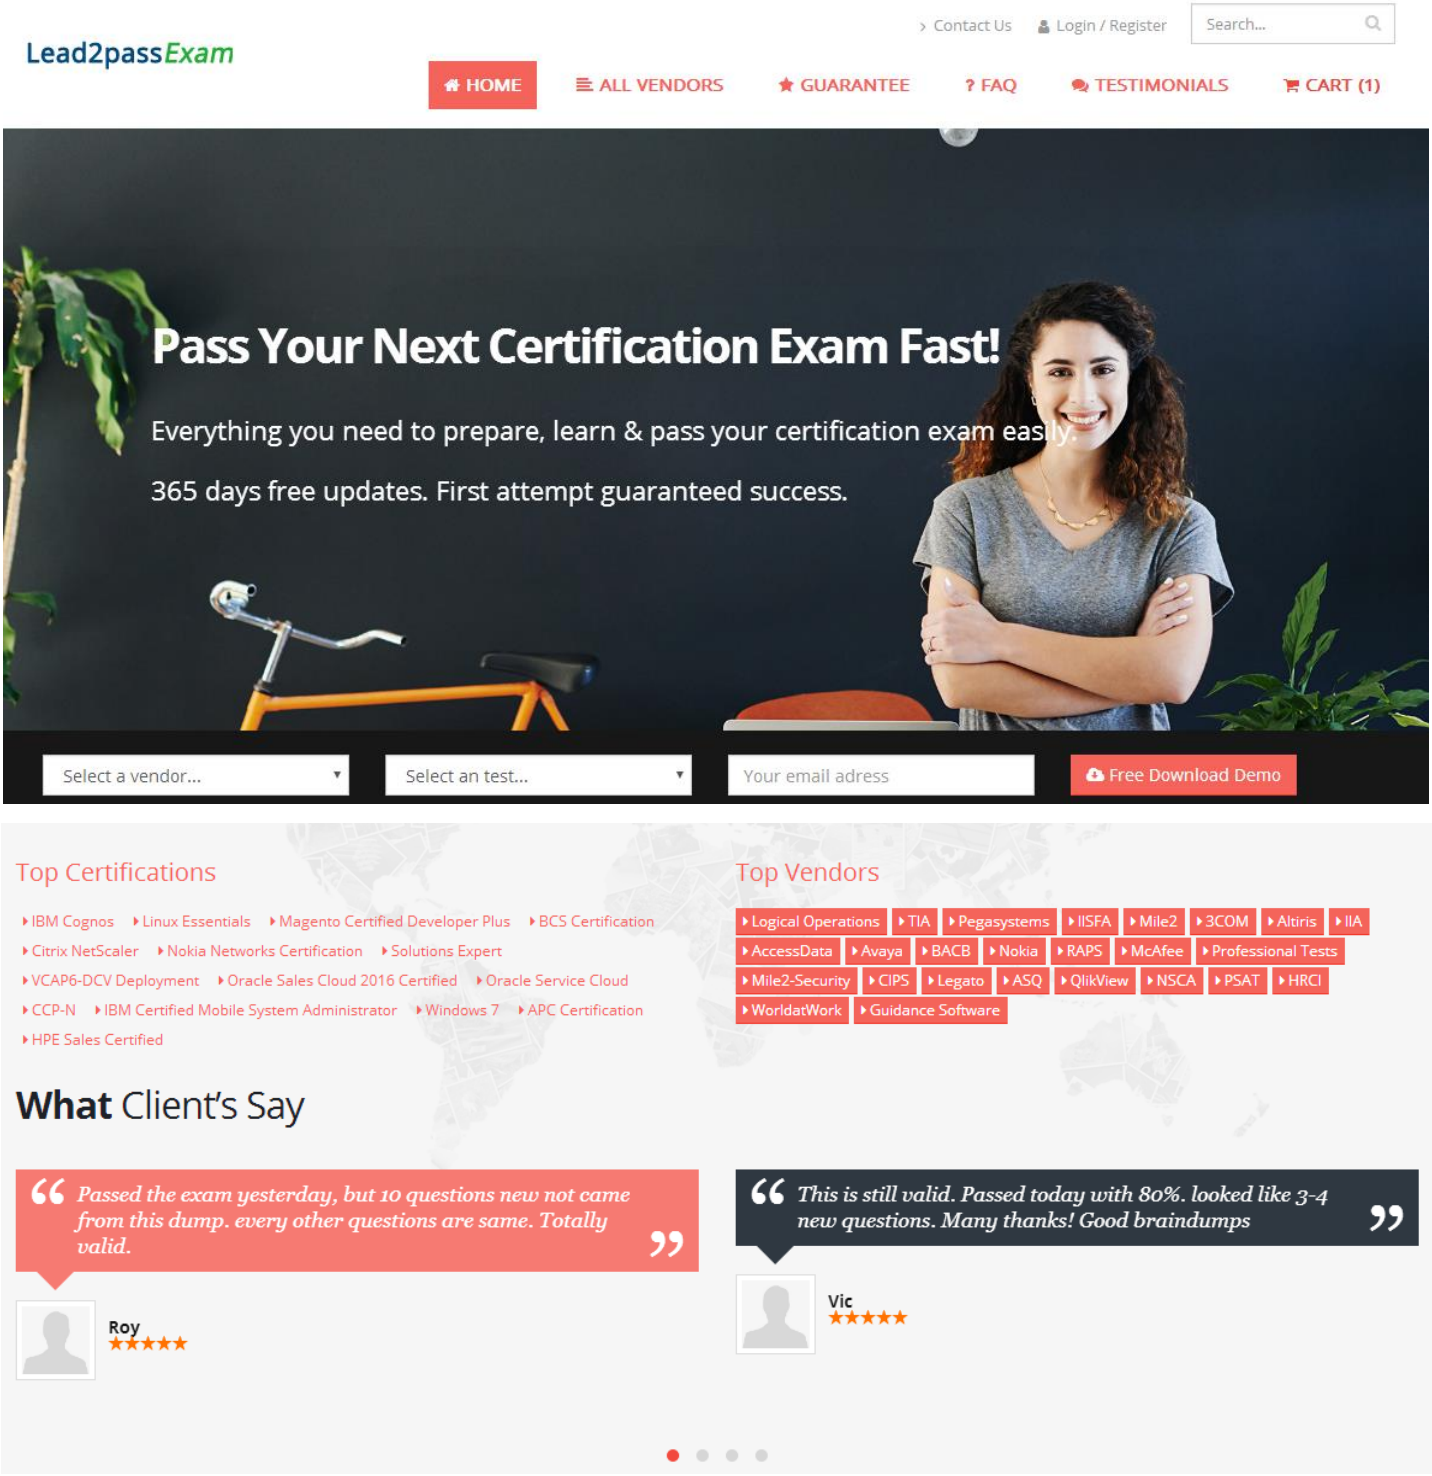

http://www.lead2passexam.com/

Available Exam Cram and Valid Dumps - Lead2Pass Exam

Exam : **H13-311_V3.5**

Title : HCIA-AI V3.5

Vendor : Huawei

Version : DEMO

NO.1 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.
Answer: A,D

NO.2 Which of the following does not belong to the process for constructing a knowledge graph?
A. Determining the target domain of the knowledge graph
B. Data acquisition
C. Creating new concepts
D. Knowledge fusion
Answer: C
Explanation:
The process of constructing a knowledge graph typically involves several key steps:
A . Determining the target domain of the knowledge graph: This defines the scope and boundaries of the information to be represented.
B . Data acquisition: Involves gathering structured and unstructured data from various sources.
D . Knowledge fusion: This step involves integrating and reconciling data from multiple sources to create a consistent and coherent knowledge graph.
Creating new concepts is not typically part of the knowledge graph construction process. Instead, knowledge graphs usually focus on extracting, integrating, and structuring existing knowledge, not creating new concepts.
HCIA AI
Reference:
AI Development Framework: Describes the steps in constructing knowledge graphs, from data acquisition to knowledge fusion and domain determination.

NO.3 HarmonyOS can provide AI capabilities for external systems only through the integrated HMS Core.
A. TRUE
B. FALSE
Answer: B
Explanation:
HarmonyOS provides AI capabilities not only through HMS Core (Huawei Mobile Services Core), but also through other system-level integrations and AI frameworks. While HMS Core is one way to offer AI functionalities, HarmonyOS also has native support for AI processing that can be accessed by external systems or applications beyond HMS Core.
Thus, the statement is false as AI capabilities are not limited solely to HMS Core in HarmonyOS.
HCIA AI
Reference:
Introduction to Huawei AI Platforms: Covers HarmonyOS and the various ways it integrates AI

capabilities into external systems.

NO.4 Nesterov is a variant of the momentum optimizer.

A. TRUE

B. FALSE

Answer: A

Explanation:

Nesterov Accelerated Gradient (NAG) is indeed a variant of the momentum optimizer. In the traditional momentum method, the gradient is used to adjust the direction based on the current momentum. Nesterov, on the other hand, anticipates the change in the momentum by calculating the gradient at a slightly altered position. This small adjustment leads to better convergence and more efficient optimization, especially in non-convex problems.

Momentum methods and their variants like Nesterov are commonly discussed in the optimization strategies for neural networks, including frameworks such as TensorFlow, which is covered in Huawei's HCIA AI courses.

HCIA AI

Reference:

Deep Learning Overview: Discussion of optimization algorithms, including gradient descent variants like Momentum and Nesterov.

AI Development Framework: Explains the use of Nesterov in deep learning frameworks such as TensorFlow and PyTorch.

NO.5 Which of the following are callback options provided by MindSpore?

A. SummaryCollector

B. TrainStep

C. ModelCheckpoint

D. LossMonitor

Answer: A,C,D

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.

NO.6 Huawei Cloud EI provides knowledge graph, OCR, machine translation, and the Celia (virtual assistant) development platform.

A. TRUE

B. FALSE

Answer: A

Explanation:

Huawei Cloud EI (Enterprise Intelligence) provides a variety of AI services and platforms, including knowledge graph, OCR (Optical Character Recognition), machine translation, and the Celia virtual

assistant development platform. These services enable businesses to integrate AI capabilities such as language processing, image recognition, and virtual assistant development into their systems.

NO.7 Which of the following is NOT a commonly used AI computing framework?
A. PyTorch
B. MindSpore
C. TensorFlow
D. OpenCV
Answer: D
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.

NO.8 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)
Answer: A,B,D
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.

NO.9 All kernels of the same convolutional layer in a convolutional neural network share a weight.
A. TRUE
B. FALSE
Answer: B
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
Reference:
Deep Learning Overview: Detailed description of CNNs, focusing on kernel operations and weight sharing mechanisms within a single kernel, but not across different kernels.

NO.10 Which of the following is the order of tensor [[0,1],[2,3]]?
A. 6
B. 3
C. 2
D. 4
Answer: C
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.

NO.11 Which of the following statements are true about decision trees?
A. The common decision tree algorithms include ID3, C4.5, and CART.
B. Quantitative indicators of purity can only be obtained by using information entropy.
C. Building a decision tree means selecting feature attributes and determining their tree structure.
D. A key step to building a decision tree involves dividing all feature attributes and comparing the purity of the division's result sets.
Answer: A,C,D
Explanation:
A . TRUE. The common decision tree algorithms include ID3, C4.5, and CART. These are the most widely used algorithms for decision tree generation.
B . FALSE. Purity in decision trees can be measured using multiple metrics, such as information gain, Gini index, and others, not just information entropy.
C . TRUE. Building a decision tree involves selecting the best features and determining their order in the tree structure to split the data effectively.
D . TRUE. One key step in decision tree generation is evaluating the purity of different splits (e.g., how well the split segregates the target variable) by comparing metrics like information gain or Gini index.
HCIA AI
Reference:
Machine Learning Overview: Covers decision tree algorithms and their use cases.
Deep Learning Overview: While this focuses on neural networks, it touches on how decision-making algorithms are used in structured data models.

NO.12 Convolutional neural networks (CNNs) cannot be used to process text data.
A. TRUE
B. FALSE
Answer: B
Explanation:
Contrary to the statement, Convolutional Neural Networks (CNNs) can indeed be used to process text data. While CNNs are most famously used for image processing, they can also be adapted for natural

language processing (NLP) tasks. In text data, CNNs can operate on word embeddings or character-level data to capture local patterns (e.g., sequences of words or characters). CNNs are used in applications such as text classification, sentiment analysis, and language modeling.
The key to CNN's application in text processing is that the convolutional layers can detect patterns in sequences, much like they detect spatial features in images. This versatility is covered in Huawei's HCIA AI platform when discussing CNN's applications beyond image data.
HCIA AI
Reference:
Deep Learning Overview: Explores the usage of CNNs in different domains, including their application in NLP tasks.
Cutting-edge AI Applications: Discusses the use of CNNs in non-traditional tasks, including text and sequential data processing.

NO.13 In a hyperparameter-based search, the hyperparameters of a model are searched based on the data on and the model's performance metrics.

A. TRUE

B. FALSE

Answer: A

Explanation:
In machine learning, hyperparameters are the parameters that govern the learning process and are not learned from the data. Hyperparameter optimization or hyperparameter tuning is a critical part of improving a model's performance. The goal of a hyperparameter-based search is to find the set of hyperparameters that maximizes the model's performance on a given dataset.
There are different techniques for hyperparameter tuning, such as grid search, random search, and more advanced methods like Bayesian optimization. The performance of the model is assessed based on evaluation metrics (like accuracy, precision, recall, etc.), and the hyperparameters are adjusted accordingly to achieve the best performance.
In Huawei's HCIA AI curriculum, hyperparameter optimization is discussed in relation to both traditional machine learning models and deep learning frameworks. The course emphasizes the importance of selecting appropriate hyperparameters and demonstrates how frameworks such as TensorFlow and Huawei's ModelArts platform can facilitate hyperparameter searches to optimize models efficiently.
HCIA AI
Reference:
AI Overview and Machine Learning Overview: Emphasize the importance of hyperparameters in model training.
Deep Learning Overview: Highlights the role of hyperparameter tuning in neural network architectures, including tuning learning rates, batch sizes, and other key parameters.
AI Development Frameworks: Discusses the use of hyperparameter search tools in platforms like TensorFlow and Huawei ModelArts.

NO.14 AI chips, also called AI accelerators, optimize matrix multiplication.

A. TRUE

B. FALSE

Answer: A

Explanation:

AI chips, also known as AI accelerators, are specialized hardware designed to enhance the performance of AI workloads, particularly for tasks like matrix multiplication, which is heavily used in machine learning and deep learning algorithms. These chips optimize operations like matrix multiplications because they are computationally intensive and central to neural network computations (e.g., in forward and backward passes).
HCIA AI
Reference:
Cutting-edge AI Applications: Discussion of AI chips and accelerators, with a focus on their role in improving computation efficiency.
Deep Learning Overview: Explains how neural network operations like matrix multiplication are optimized in AI hardware.

NO.15 "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.

Answer: A
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.

NO.16 Which of the following activation functions may cause the vanishing gradient problem?

A. Softplus

B. ReLU

C. Sigmoid

D. Tanh

Answer: C,D
Explanation:
Both Sigmoid and Tanh activation functions can cause the vanishing gradient problem. This issue occurs because these functions squash their inputs into a very small range, leading to very small gradients during backpropagation, which slows down learning. In deep neural networks, this can prevent the weights from updating effectively, causing the training process to stall.
Sigmoid: Outputs values between 0 and 1. For large positive or negative inputs, the gradient becomes very small.
Tanh: Outputs values between -1 and 1. While it has a broader range than Sigmoid, it still suffers from vanishing gradients for larger input values.
ReLU, on the other hand, does not suffer from the vanishing gradient problem since it outputs the input directly if positive, allowing gradients to pass through. However, Softplus is also less prone to

this problem compared to Sigmoid and Tanh.
HCIA AI
Reference:
Deep Learning Overview: Explains the vanishing gradient problem in deep networks, especially when using Sigmoid and Tanh activation functions.
AI Development Framework: Covers the use of ReLU to address the vanishing gradient issue and its prevalence in modern neural networks.

NO.17 Fill in blanks
The general process of building a project using machine learning involves the following steps: split data, _____ the model, deploy the model the model, and fine-tune the model.
Answer: train

NO.18 The concept of "artificial intelligence" was first proposed in the year of:
A. 1950
B. 1956
C. 1960
D. 1965
Answer: B
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.

NO.19 When learning the MindSpore framework, John learns how to use callbacks and wants to use it for AI model training. For which of the following scenarios can John use the callback?
A. Early stopping
B. Adjusting an activation function
C. Saving model parameters
D. Monitoring loss values during training
Answer: A,C,D
Explanation:
In MindSpore, callbacks can be used in various scenarios such as:
Early stopping: To stop training when the performance plateaus or certain criteria are met.
Saving model parameters: To save checkpoints during or after training using the ModelCheckpoint callback.
Monitoring loss values: To keep track of loss values during training using LossMonitor, allowing interventions if necessary.
Adjusting the activation function is not a typical use case for callbacks, as activation functions are usually set during model definition.

NO.20 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

Answer: A

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.

NO.21 Huawei Cloud ModelArts provides ModelBox for device-edge-cloud joint development. Which of the following are its optimization policies?

A. Hardware affinity

B. Operator optimization

C. Automatic segmentation of operators

D. Model replication

Answer: A,B,C

Explanation:

Huawei Cloud ModelArts provides ModelBox, a tool for device-edge-cloud joint development, enabling efficient deployment across multiple environments. Some of its key optimization policies include:

Hardware affinity: Ensures that the models are optimized to run efficiently on the target hardware.

Operator optimization: Improves the performance of AI operators for better model execution.

Automatic segmentation of operators: Automatically segments operators for optimized distribution across devices, edges, and clouds.

Model replication is not an optimization policy offered by ModelBox.

NO.22 The derivative of the Rectified Linear Unit (ReLU) activation function in the positive interval is always:

A. 0

B. 0.5

C. 1

D. Variable

Answer: C

Explanation:

The Rectified Linear Unit (ReLU) activation function is defined as f(x)=max(0,x)f(x) = \max(0, x)f(x)=max(0,x). In the positive interval, where x>0x > 0x>0, the derivative of ReLU is always 1. This makes ReLU popular for deep learning networks because it helps avoid the vanishing gradient problem during backpropagation, ensuring efficient gradient flow.

NO.23 When feature engineering is complete, which of the following is not a step in the decision tree building process?

A. Decision tree generation
B. Pruning
C. Feature selection
D. Data cleansing
Answer: D
Explanation:
When building a decision tree, the steps generally involve:
Decision tree generation: This is the process where the model iteratively splits the data based on feature values to form branches.
Pruning: This step occurs post-generation, where unnecessary branches are removed to reduce overfitting and enhance generalization.
Feature selection: This is part of decision tree construction, where relevant features are selected at each node to determine how the tree branches.
Data cleansing, on the other hand, is a preprocessing step carried out before any model training begins. It involves handling missing or erroneous data to improve the quality of the dataset but is not part of the decision tree building process itself.
HCIA AI
Reference:
Machine Learning Overview: Includes a discussion on decision tree algorithms and the process of building decision trees.
AI Development Framework: Highlights the steps for building machine learning models, separating data preprocessing (e.g., data cleansing) from model building steps.

NO.24 Which of the following are subfields of AI?
A. Backpropagation algorithm
B. Expert system
C. Smart finance
D. Computer vision
Answer: B,D
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.

NO.25 Which of the following algorithms presents the most chaotic landscape on the loss surface?
A. Stochastic gradient descent
B. MGD
C. MBGD
D. BGD
Answer: A

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.

NO.26 In MindSpore, the basic unit of the neural network is nn.Cell.
A. TRUE
B. FALSE
Answer: A
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.

NO.27 Which of the following is the activation function used in the hidden layers of the standard recurrent neural network (RNN) structure?
A. ReLU
B. Softmax
C. Tanh
D. Sigmoid
Answer: C
Explanation:
In standard Recurrent Neural Networks (RNNs), the Tanh activation function is commonly used in the hidden layers. The Tanh function squashes input values to a range between -1 and 1, allowing the network to learn complex patterns over time by transforming the input data into non-linear patterns. While other activation functions like Sigmoid can be used, Tanh is preferred in many RNNs for its wider range. ReLU is generally used in feed-forward networks, and Softmax is often applied in the output layer for classification problems.
HCIA AI
Reference:
Deep Learning Overview: Describes the architecture of RNNs, highlighting the use of Tanh as the standard activation function.
AI Development Framework: Discusses the various activation functions used across different neural network architectures.

NO.28 The mean squared error (MSE) loss function cannot be used for classification problems.
A. TRUE
B. FALSE
Answer: A
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.

NO.29 Single-layer perceptrons and logistic regression are linear classifiers that can only process linearly separable data.

A. TRUE

B. FALSE

Answer: A

Explanation:
Both single-layer perceptrons and logistic regression are linear classifiers, meaning they are capable of separating data that is linearly separable. However, they cannot effectively model non-linear relationships in the data. For more complex, non-linearly separable data, multi-layer neural networks or other non-linear classifiers are required.

NO.30 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

Answer: A

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
Reference:
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.

NO.31 Sigmoid, tanh, and softsign activation functions cannot avoid vanishing gradient problems when the network is deep.

A. TRUE

B. FALSE

Answer: A

Explanation:
Activation functions like Sigmoid, tanh, and softsign suffer from the vanishing gradient problem when used in deep networks. This happens because, in these functions, gradients become very small as the input moves away from the origin (either positively or negatively). As a result, the weights of the

earlier layers in the network receive very small updates, hindering the learning process in deep networks. This is one reason why activation functions like ReLU, which avoid this issue, are often preferred in deep learning.

NO.32 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

Answer: C

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.

NO.33 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

Answer: A

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

Reference:

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.

NO.34 Correctly connect the layers in the architecture of an Ascend AI Processor.

Answer:

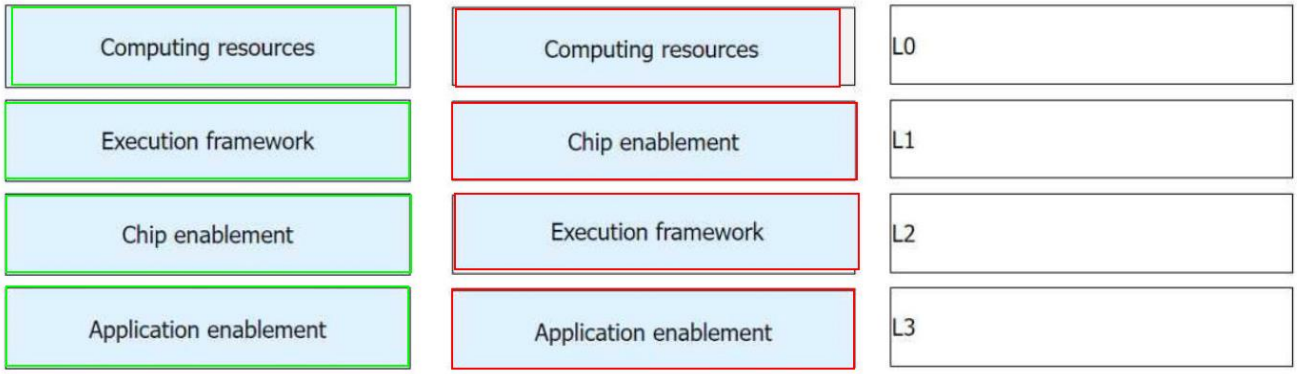

NO.35 When you use MindSpore to execute the following code, which of the following is the output?
python
Copy code
x = Tensor(np.array([[1, 2], [3, 4]]), dtype.int32)
x.dtype

A. 2

B. mindspore.int32

C. 4

D. (2,2)

Answer: B

Explanation:
In MindSpore, when you define a tensor using Tensor(np.array([[1, 2], [3, 4]]), dtype.int32), the dtype attribute returns the data type of the tensor, which in this case is mindspore.int32. This specifies the type of elements in the tensor.