Free Questions for AIP-210

Shared by Bullock on 12-12-2023

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Which of the following is the primary purpose of hyperparameter optimization?

Options:

A- Controls the learning process of a given algorithm
B- Makes models easier to explain to business stakeholders
C- Improves model interpretability
D- Increases recall over precision

Answer:

A

Explanation:

Hyperparameter optimization is the process of finding the optimal values for hyperparameters that control the learning process of a given algorithm. Hyperparameters are parameters that are not learned by the algorithm but are set by the user before training. Hyperparameters can affect the performance and behavior of the algorithm, such as its speed, accuracy, complexity, or generalization. Hyperparameter optimization can help improve the efficiency and effectiveness of the algorithm by tuning its hyperparameters to achieve the best results.

Question 2

Question Type: MultipleChoice

Which of the following items should be included in a handover to the end user to enable them to use and run a trained model on their own system? (Select three.)

Options:

A- Information on the folder structure in your local machine
B- Intermediate data files
C- Link to a GitHub repository of the codebase
D- README document

E- Sample input and output data files

Answer:

C, D, E

Explanation:

A handover is the process of transferring the ownership and responsibility of an ML system from one party to another, such as from the developers to the end users. A handover should include all the necessary information and resources that enable the end users to use and run a trained model on their own system. Some of the items that should be included in a handover are:

Link to a GitHub repository of the codebase: A GitHub repository is an online platform that hosts the source code and version control of an ML system. A link to a GitHub repository can provide the end users with access to the latest and most updated version of the codebase, as well as the history and documentation of the changes made to the code.

README document: A README document is a text file that provides an overview and instructions for an ML system. A README document can include information such as the purpose, features, requirements, installation, usage, testing, troubleshooting, and license of the system.

Sample input and output data files: Sample input and output data files are data files that contain examples of valid inputs and expected outputs for an ML system. Sample input and output data files can help the end users understand how to use and run the system, as well as verify its functionality and performance.

Question 3

Question Type: MultipleChoice

An AI practitioner incorporates risk considerations into a deployment plan and decides to log and store historical predictions for potential, future access requests.

Which ethical principle is this an example of?

Options:

A- Fairness
B- Privacy
C- Safety
D- Transparency

Answer:

D

Explanation:

Transparency is an ethical principle that describes the degree to which an AI system can provide clear and understandable information about its inputs, outputs, processes, and decisions. Transparency can help increase trust and confidence among users and stakeholders, as well as enable accountability and responsibility for the system's actions and outcomes. Logging and storing historical predictions for potential, future access requests is an example of transparency, as it can help provide evidence and explanation for the system's recommendations, as well as facilitate auditing and feedback.

Question 4

Question Type: MultipleChoice

The following confusion matrix is produced when a classifier is used to predict labels on a test dataset. How precise is the classifier?



Options:

A- 48/(48+37)
B- 37/(37+8)
C- 37/(37+7)
D- (48+37)/100

Answer:

B

Explanation:

Precision is a measure of how well a classifier can avoid false positives (incorrectly predicted positive cases). Precision is calculated by dividing the number of true positives (correctly predicted positive cases) by the number of predicted positive cases (true positives and false positives). In this confusion matrix, the true positives are 37 and the false positives are 8, so the

precision is 37/(37+8) = 0.822.

Question 5

Question Type: MultipleChoice

Which of the following best describes distributed artificial intelligence?

Options:

A- It does not require hyperparemeter tuning because the distributed nature accounts for the bias.
B- It intelligently pre-distributes the weight of starting a neural network.
C- It relies on a distributed system that performs robust computations across a network of unreliable nodes.
D- It uses a centralized system to speak to decentralized nodes.

Answer:

C

Explanation:

Distributed artificial intelligence (DAI) is a subfield of artificial intelligence that studies how multiple intelligent agents can coordinate and cooperate to achieve a common goal or solve a complex problem. DAI relies on a distributed system that performs robust computations across a network of unreliable nodes, such as sensors, robots, or humans. DAI can handle large-scale, dynamic, and uncertain environments that are beyond the capabilities of a single agent.
Reference: [Distributed artificial intelligence - Wikipedia], [Distributed Artificial Intelligence: An Overview]

Question 6

Question Type: MultipleChoice

Which of the following equations best represent an LI norm?

Options:

A- |x| + |y|
B- |x|+|y|^2
C- |x|-|y|
D- |x|^2+|y|^2

Answer:

A

Explanation:

An L1 norm is a measure of distance or magnitude that is defined as the sum of the absolute values of the components of a vector. For example, if x and y are two components of a vector, then the L1 norm of that vector is |x| + |y|. The L1 norm is also known as the Manhattan distance or the taxicab distance, as it represents the shortest path between two points in a grid-like city.

Question 7

Question Type: MultipleChoice

Which of the following is NOT an activation function?

Options:

A- Additive
B- Hyperbolic tangent
C- ReLU
D- Sigmoid

Answer:

A

Explanation:

An activation function is a function that determines the output of a neuron in a neural network based on its input. An activation function can introduce non-linearity into a neural network, which allows it to model complex and non-linear relationships between inputs and outputs. Some of the

common activation functions are:

Sigmoid: A sigmoid function is a function that maps any real value to a value between 0 and 1. It has an S-shaped curve and is often used for binary classification or probability estimation.

Hyperbolic tangent: A hyperbolic tangent function is a function that maps any real value to a value between -1 and 1. It has a similar shape to the sigmoid function but is symmetric around the origin. It is often used for regression or classification problems.

ReLU: A ReLU (rectified linear unit) function is a function that maps any negative value to 0 and any positive value to itself. It has a piecewise linear shape and is often used for hidden layers in deep neural networks.

Additive is not an activation function, but rather a term that describes a property of some functions. Additive functions are functions that satisfy the condition f(x+y) = f(x) + f(y) for any x and y. Additive functions are linear functions, which means they have a constant slope and do not introduce non-linearity.

Question 8

Question Type: MultipleChoice

Workflow design patterns for the machine learning pipelines:

Options:

A- Aim to explain how the machine learning model works.
B- Represent a pipeline with directed acyclic graph (DAG).
C- Seek to simplify the management of machine learning features.
D- Separate inputs from features.

Answer:

B

Explanation:

Workflow design patterns for machine learning pipelines are common solutions to recurring problems in building and managing machine learning workflows. One of these patterns is to represent a pipeline with a directed acyclic graph (DAG), which is a graph that consists of nodes and edges, where each node represents a step or task in the pipeline, and each edge represents a dependency or order between the tasks. A DAG has no cycles, meaning there is no way to start

at one node and return to it by following the edges. A DAG can help visualize and organize the pipeline, as well as facilitate parallel execution, fault tolerance, and reproducibility.

Question 9

Question Type: MultipleChoice

You create a prediction model with 96% accuracy. While the model's true positive rate (TPR) is performing well at 99%, the true negative rate (TNR) is only 50%. Your supervisor tells you that the TNR needs to be higher, even if it decreases the TPR. Upon further inspection, you notice that the vast majority of your data is truly positive.

What method could help address your issue?

Options:

A- Normalization
B- Oversampling
C- Principal components analysis
D- Quality filtering

Answer:

B

Explanation:

Oversampling is a method that can help address the issue of imbalanced data, which is when one class is much more frequent than the other in the dataset. This can cause the model to be biased towards the majority class and have a low true negative rate. Oversampling involves creating synthetic samples of the minority class or replicating existing samples to balance the class distribution. This can help the model learn more from the minority class and improve the true negative rate. Reference: [Handling imbalanced datasets in machine learning], [Oversampling and undersampling in data analysis - Wikipedia]

To Get Premium Files for AIP-210 Visit
https://www.p2pexams.com/products/aip-210

For More Free Questions Visit
https://www.p2pexams.com/certnexus/pdf/aip-210

