

Free Questions for CT-AI

Shared by Savage on 04-10-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Which ONE of the following statements is a CORRECT adversarial example in the context of machine learning systems that are working on image classifiers.

SELECT ONE OPTION

Options:

- A-** Black box attacks based on adversarial examples create an exact duplicate model of the original.
- B-** These attack examples cause a model to predict the correct class with slightly less accuracy even though they look like the original image.
- C-** These attacks can't be prevented by retraining the model with these examples augmented to the training data.
- D-** These examples are model specific and are not likely to cause another model trained on same task to fail.

Answer:

D

Explanation:

A . Black box attacks based on adversarial examples create an exact duplicate model of the original.

Black box attacks do not create an exact duplicate model. Instead, they exploit the model by querying it and using the outputs to craft adversarial examples without knowledge of the internal workings.

B . These attack examples cause a model to predict the correct class with slightly less accuracy even though they look like the original image.

Adversarial examples typically cause the model to predict the incorrect class rather than just reducing accuracy. These examples are designed to be visually indistinguishable from the original image but lead to incorrect classifications.

C . These attacks can't be prevented by retraining the model with these examples augmented to the training data.

This statement is incorrect because retraining the model with adversarial examples included in the training data can help the model learn to resist such attacks, a technique known as adversarial training.

D . These examples are model specific and are not likely to cause another model trained on the same task to fail.

Adversarial examples are often model-specific, meaning that they exploit the specific weaknesses of a particular model. While some adversarial examples might transfer between models, many are tailored to the specific model they were generated for and may not affect other models trained on the same task.

Therefore, the correct answer is D because adversarial examples are typically model-specific and may not cause another model trained on the same task to fail.

Question 2

Question Type: MultipleChoice

Which ONE of the following hardware is MOST suitable for implementing AI when using ML?

SELECT ONE OPTION

Options:

- A- 64-bit CPUs.
- B- Hardware supporting fast matrix multiplication.
- C- High powered CPUs.
- D- Hardware supporting high precision floating point operations.

Answer:

B

Explanation:

A . 64-bit CPUs.

While 64-bit CPUs are essential for handling large amounts of memory and performing complex computations, they are not specifically optimized for the types of operations commonly used in machine learning.

B . Hardware supporting fast matrix multiplication.

Matrix multiplication is a fundamental operation in many machine learning algorithms, especially in neural networks and deep learning. Hardware optimized for fast matrix multiplication, such as GPUs (Graphics Processing Units), is most suitable for implementing AI and ML because it can handle the parallel processing required for these operations efficiently.

C . High powered CPUs.

High powered CPUs are beneficial for general-purpose computing tasks and some aspects of ML, but they are not as efficient as specialized hardware like GPUs for matrix multiplication and other ML-specific tasks.

D . Hardware supporting high precision floating point operations.

High precision floating point operations are important for scientific computing and some specific AI tasks, but for many ML applications, fast matrix multiplication is more critical than high precision alone.

Therefore, the correct answer is B because hardware supporting fast matrix multiplication, such as GPUs, is most suitable for the parallel processing requirements of machine learning.

Question 3

Question Type: MultipleChoice

In a conference on artificial intelligence (AI), a speaker made the statement, "The current implementation of AI using models which do NOT change by themselves is NOT true AI". Based on your understanding of AI, is this above statement CORRECT or INCORRECT and why?

SELECT ONE OPTION

Options:

- A-** This statement is incorrect. Current AI is true AI and there is no reason to believe that this fact will change over time.
- B-** This statement is correct. In general, what is considered AI today may change over time.
- C-** This statement is incorrect. What is considered AI today will continue to be AI even as technology evolves and changes.
- D-** This statement is correct. In general, today the term AI is utilized incorrectly.

Answer:

B

Explanation:

A. This statement is incorrect. Current AI is true AI and there is no reason to believe that this fact will change over time.

AI is an evolving field, and the definition of what constitutes AI can change as technology advances.

B. This statement is correct. In general, what is considered AI today may change over time.

The term AI is dynamic and has evolved over the years. What is considered AI today might be viewed as standard computing in the future. Historically, as technologies become mainstream, they often cease to be considered 'AI'.

C. This statement is incorrect. What is considered AI today will continue to be AI even as technology evolves and changes.

This perspective does not account for the historical evolution of the definition of AI . As new technologies emerge, the boundaries of AI shift.

D. This statement is correct. In general, today the term AI is utilized incorrectly.

While some may argue this, it is not a universal truth. The term AI encompasses a broad range of technologies and applications, and its usage is generally consistent with current technological capabilities.

Question 4

Question Type: MultipleChoice

Which ONE of the following activities is MOST relevant when addressing the scenario where you have more than the required amount of data available for the training?

SELECT ONE OPTION

Options:

- A- Feature selection
- B- Data sampling
- C- Data labeling
- D- Data augmentation

Answer:

B

Explanation:

A . Feature selection

Feature selection is the process of selecting the most relevant features from the data. While important, it is not directly about handling excess data.

B . Data sampling

Data sampling involves selecting a representative subset of the data for training. When there is more data than needed, sampling can be used to create a manageable dataset that maintains the statistical properties of the full dataset.

C . Data labeling

Data labeling involves annotating data for supervised learning. It is necessary for training models but does not address the issue of having excess data.

D . Data augmentation

Data augmentation is used to increase the size of the training dataset by creating modified versions of existing data. It is useful when there is insufficient data, not when there is excess data.

Therefore, the correct answer is B because data sampling is the most relevant activity when dealing with an excess amount of data for training.

Question 5

Question Type: MultipleChoice

Which ONE of the following options describes the LEAST LIKELY usage of AI for detection of GUI changes due to changes in test objects?

SELECT ONE OPTION

Options:

- A- Using a pixel comparison of the GUI before and after the change to check the differences.
- B- Using a computer vision to compare the GUI before and after the test object changes.
- C- Using a vision-based detection of the GUI layout changes before and after test object changes.
- D- Using a ML-based classifier to flag if changes in GUI are to be flagged for humans.

Answer:

A

Explanation:

A. Using a pixel comparison of the GUI before and after the change to check the differences.

Pixel comparison is a traditional method and does not involve AI . It compares images at the pixel level, which can be effective but is not an intelligent approach. It is not considered an AI usage and is the least likely usage of AI for detecting GUI changes.

B. Using computer vision to compare the GUI before and after the test object changes.

Computer vision involves using AI techniques to interpret and process images. It is a likely usage of AI for detecting changes in the GUI .

C. Using vision-based detection of the GUI layout changes before and after test object changes.

Vision-based detection is another AI technique where the layout and structure of the GUI are analyzed to detect changes. This is a typical application of AI .

D. Using a ML-based classifier to flag if changes in GUI are to be flagged for humans.

An ML-based classifier can intelligently determine significant changes and decide if they need human review, which is a sophisticated AI application.

Question 6

Question Type: MultipleChoice

"Splendid Healthcare" has started developing a cancer detection system based on ML. The type of cancer they plan on detecting has 2% prevalence rate in the population of a particular geography. It is required that the model performs well for both normal and cancer patients.

Which ONE of the following combinations requires MAXIMIZATION?

SELECT ONE OPTION

Options:

A- Maximize precision and accuracy

- B-** Maximize accuracy and recall
- C-** Maximize recall and precision
- D-** Maximize specificity number of classes

Answer:

C

Explanation:

Prevalence Rate and Model Performance:

The cancer detection system being developed by 'Splendid Healthcare' needs to account for the fact that the type of cancer has a 2% prevalence rate in the population. This indicates that the dataset is highly imbalanced with far fewer positive (cancer) cases compared to negative (normal) cases.

Importance of Recall:

Recall, also known as sensitivity or true positive rate, measures the proportion of actual positive cases that are correctly identified by the model. In medical diagnosis, especially cancer detection, recall is critical because missing a positive case (false negative) could have severe consequences for the patient. Therefore, maximizing recall ensures that most, if not all, cancer cases are detected.

Importance of Precision:

Precision measures the proportion of predicted positive cases that are actually positive. High precision reduces the number of false positives, meaning fewer people will be incorrectly diagnosed with cancer. This is also important to avoid unnecessary anxiety and further invasive testing for those who do not have the disease.

Balancing Recall and Precision:

In scenarios where both false negatives and false positives have significant consequences, it is crucial to balance recall and precision. This balance ensures that the model is not only good at detecting positive cases but also accurate in its predictions, reducing both types of errors.

Accuracy and Specificity:

While accuracy (the proportion of total correct predictions) is important, it can be misleading in imbalanced datasets. In this case, high accuracy could simply result from the model predicting the majority class (normal) correctly. Specificity (true negative rate) is also important, but for a cancer detection system, recall and precision take precedence to ensure positive cases are correctly and accurately identified.

Conclusion:

Therefore, for a cancer detection system with a low prevalence rate, maximizing both recall and precision is crucial to ensure effective and accurate detection of cancer cases.

Question 7

Question Type: MultipleChoice

A ML engineer is trying to determine the correctness of the new open-source implementation *X", of a supervised regression algorithm implementation. R-Square is one of the functional performance metrics used to determine the quality of the model.

Which ONE of the following would be an APPROPRIATE strategy to achieve this goal?

SELECT ONE OPTION

Options:

- A-** Add 10% of the rows randomly and create another model and compare the R-Square scores of both the model.
- B-** Train various models by changing the order of input features and verify that the R-Square score of these models vary significantly.
- C-** Compare the R-Square score of the model obtained using two different implementations that utilize two different programming languages while using the same algorithm and the same training and testing data.
- D-** Drop 10% of the rows randomly and create another model and compare the R-Square scores of both the models.

Answer:

C

Explanation:

A . Add 10% of the rows randomly and create another model and compare the R-Square scores of both the models.

Adding more data to the training set can affect the R-Square score, but it does not directly verify the correctness of the implementation.

B . Train various models by changing the order of input features and verify that the R-Square score of these models vary significantly.

Changing the order of input features should not significantly affect the R-Square score if the implementation is correct, but this approach is more about testing model robustness rather than correctness of the implementation.

C . Compare the R-Square score of the model obtained using two different implementations that utilize two different programming languages while using the same algorithm and the same training and testing data.

This approach directly compares the performance of two implementations of the same algorithm. If both implementations produce similar R-Square scores on the same training and testing data, it suggests that the new implementation 'X' is correct.

D . Drop 10% of the rows randomly and create another model and compare the R-Square scores of both the models.

Dropping data can lead to variations in the R-Square score but does not directly verify the correctness of the implementation.

Therefore, option C is the most appropriate strategy because it directly compares the performance of the new implementation 'X' with another implementation using the same algorithm and datasets, which helps in verifying the correctness of the implementation.

Question 8

Question Type: MultipleChoice

An image classification system is being trained for classifying faces of humans. The distribution of the data is 70% ethnicity A and 30% for ethnicities B, C and D. Based ONLY on the above information, which of the following options BEST describes the situation of this image classification system?

SELECT ONE OPTION

Options:

- A- This is an example of expert system bias.
- B- This is an example of sample bias.
- C- This is an example of hyperparameter bias.
- D- This is an example of algorithmic bias.

Answer:

B

Explanation:

A . This is an example of expert system bias.

Expert system bias refers to bias introduced by the rules or logic defined by experts in the system, not by the data distribution.

B . This is an example of sample bias.

Sample bias occurs when the training data is not representative of the overall population that the model will encounter in practice. In this case, the over-representation of ethnicity A (70%) compared to B, C, and D (30%) creates a sample bias, as the model may become biased towards better performance on ethnicity A.

C . This is an example of hyperparameter bias.

Hyperparameter bias relates to the settings and configurations used during the training process, not the data distribution itself.

D . This is an example of algorithmic bias.

Algorithmic bias refers to biases introduced by the algorithmic processes and decision-making rules, not directly by the distribution of training data.

Based on the provided information, option B (sample bias) best describes the situation because the training data is skewed towards ethnicity A, potentially leading to biased model performance.

Question 9

Question Type: MultipleChoice

In the near future, technology will have evolved, and AI will be able to learn multiple tasks by itself without needing to be retrained, allowing it to operate even in new environments. The cognitive abilities of AI are similar to a child of 1-2 years.'

In the above quote, which ONE of the following options is the correct name of this type of AI?

SELECT ONE OPTION

Options:

- A- Technological singularity
- B- Narrow AI
- C- Super AI
- D- General AI

Answer:

D

Explanation:

A. Technological singularity

Technological singularity refers to a hypothetical point in the future when AI surpasses human intelligence and can continuously improve itself without human intervention. This scenario involves capabilities far beyond those described in the question.

B. Narrow AI

Narrow AI, also known as weak AI, is designed to perform a specific task or a narrow range of tasks. It does not have general cognitive abilities and cannot learn multiple tasks by itself without retraining.

C. Super AI

Super AI refers to an AI that surpasses human intelligence and capabilities across all fields. This is an advanced concept and not aligned with the description of having cognitive abilities similar to a young child.

D. General AI

General AI, or strong AI, has the ability to understand, learn, and apply knowledge across a wide range of tasks, similar to human cognitive abilities. It aligns with the description of AI that can learn multiple tasks and operate in new environments without needing retraining.

Question 10

Question Type: MultipleChoice

"AllerEgo" is a product that uses self-learning to predict the behavior of a pilot under combat situation for a variety of terrains and enemy aircraft formations. Post training the model was exposed to the real-

world data and the model was found to be behaving poorly. A lot of data quality tests had been performed on the data to bring it into a shape fit for training and testing.

Which ONE of the following options is least likely to describes the possible reason for the fall in the performance, especially when considering the self-learning nature of the AI system?

SELECT ONE OPTION

The difficulty of defining criteria for improvement before the model can be accepted.

The fast pace of change did not allow sufficient time for testing.

The unknown nature and insufficient specification of the operating environment might have caused the poor performance.

There was an algorithmic bias in the AI system.

Options:

A- The difficulty of defining criteria for improvement before the model can be accepted.

Defining criteria for improvement is a challenge in the acceptance of AI models, but it is not directly related to the performance drop in real-world scenarios. It relates more to the evaluation and deployment phase rather than affecting the model's real-time performance post-deployment.

B- The fast pace of change did not allow sufficient time for testing.

This can significantly affect the model's performance. If the system is self-learning, it needs to adapt quickly, and insufficient testing time can lead to incomplete learning and poor performance.

C- The unknown nature and insufficient specification of the operating environment might have caused the poor performance.

This is highly likely to affect performance. Self-learning AI systems require detailed specifications of the operating environment to adapt and learn effectively. If the environment is insufficiently specified, the model may fail to perform accurately in real-world scenarios.

D- There was an algorithmic bias in the AI system.

Algorithmic bias can significantly impact the performance of AI systems. If the model has biases, it will not perform well across different scenarios and data distributions.

Given the context of the self-learning nature and the need for real-time adaptability, option A is least likely to describe the fall in performance because it deals with acceptance criteria rather than real-time performance issues.

Answer:

A

Question 11

Question Type: MultipleChoice

Which of the following is THE LEAST appropriate tests to be performed for testing a feature related to autonomy?

SELECT ONE OPTION

Options:

A- Test for human handover to give rest to the system.

- B-** Test for human handover when it should actually not be relinquishing control.
- C-** Test for human handover requiring mandatory relinquishing control.
- D-** Test for human handover after a given time interval.

Answer:

B

Explanation:

Testing Autonomy: Testing for human handover when it should not be relinquishing control is the least appropriate because it contradicts the very definition of autonomous systems. The other tests are relevant to ensuring smooth operation and transitions between human and AI control.

Reference: ISTQB_CT-AI_Syllabus_v1.0, Sections on Testing Autonomous AI-Based Systems and Testing for Human-AI Interaction.

To Get Premium Files for CT-AI Visit

<https://www.p2pexams.com/products/ct-ai>

For More Free Questions Visit

<https://www.p2pexams.com/isqi/pdf/ct-ai>

