AIF-C01 Q&As

Amazon AWS Certified AI Practitioner

Pass Amazon AIF-C01 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.leads4pass.com/aif-c01.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon Official Exam Center

QUESTION 1

An AI practitioner has built a deep learning model to classify the types of materials in images. The AI practitioner now wants to measure the model performance.

Which metric will help the AI practitioner evaluate the performance of the model?

A. Confusion matrix

B. Correlation matrix

C. R2 score

D. Mean squared error (MSE)

Correct Answer: A

A confusion matrix is the correct metric for evaluating the performance of a classification model, such as the deep learning model built to classify types of materials in images.

QUESTION 2

A company has installed a security camera. The company uses an ML model to evaluate the security camera footage for potential thefts. The company has discovered that the model disproportionately flags people who are members of a specific ethnic group. Which type of bias is affecting the model output?

A. Measurement bias

B. Sampling bias

C. Observer bias

D. Confirmation bias

Correct Answer: B

Sampling bias is the correct type of bias affecting the model output when it disproportionately flags people from a specific ethnic group.

QUESTION 3

A financial institution is using Amazon Bedrock to develop an AI application. The application is hosted in a VPC. To meet regulatory compliance standards, the VPC is not allowed access to any internet traffic.

Which AWS service or feature will meet these requirements?

A. AWS PrivateLink

B. Amazon Macie

C. Amazon CloudFront

D. Internet gateway

Correct Answer: A

AWS PrivateLink enables private connectivity between VPCs and AWS services without exposing traffic to the public internet. This feature is critical for meeting regulatory compliance standards that require isolation from public internet traffic.

Option A (Correct): "AWS PrivateLink":This is the correct answer because it allows secure access to Amazon Bedrock and other AWS services from a VPC without internet access, ensuring compliance with regulatory standards. Option

B:"Amazon Macie" is incorrect because it is a security service for data classification and protection, not for managing private network traffic. Option C:"Amazon CloudFront" is incorrect because it is a content delivery network service and does

not provide private network connectivity. Option D:"Internet gateway" is incorrect as it enables internet access, which violates the VPC\\'s no-internet-traffic policy.

AWS AI Practitioner References:

AWS PrivateLink Documentation:AWS highlights PrivateLink as a solution for connecting VPCs to AWS services privately, which is essential for organizations with strict regulatory requirements.

QUESTION 4

A company is using a pre-trained large language model (LLM) to build a chatbot for product recommendations. The company needs the LLM outputs to be short and written in a specific language.

Which solution will align the LLM response quality with the company\\'s expectations?

A. Adjust the prompt.

B. Choose an LLM of a different size.

C. Increase the temperature.

D. Increase the Top K value.

Correct Answer: A

Adjusting the prompt is the correct solution to align the LLM outputs with the company\\'s expectations for short, specific language responses.

QUESTION 5

A security company is using Amazon Bedrock to run foundation models (FMs). The company wants to ensure that only authorized users invoke the models. The company needs to identify any unauthorized access attempts to set appropriate AWS Identity and Access Management (IAM) policies and roles for future iterations of the FMs.

Which AWS service should the company use to identify unauthorized users that are trying to access Amazon Bedrock?

A. AWS Audit Manager

B. AWS CloudTrail

C. Amazon Fraud Detector

D. AWS Trusted Advisor

Correct Answer: B

AWS CloudTrail is a service that enables governance, compliance, and operational and risk auditing of your AWS account. It tracks API calls and identifies unauthorized access attempts to AWS resources, including Amazon Bedrock.

QUESTION 6

A company is using domain-specific models. The company wants to avoid creating new models from the beginning. The company instead wants to adapt pre-trained models to create models for new, related tasks.

Which ML strategy meets these requirements?

A. Increase the number of epochs.

B. Use transfer learning.

C. Decrease the number of epochs.

D. Use unsupervised learning.

Correct Answer: B

Transfer learning is the correct strategy for adapting pre-trained models for new, related tasks without creating models from scratch.

QUESTION 7

An AI practitioner is using an Amazon Bedrock base model to summarize session chats from the customer service department. The AI practitioner wants to store invocation logs to monitor model input and output data.

Which strategy should the AI practitioner use?

A. Configure AWS CloudTrail as the logs destination for the model.

B. Enable invocation logging in Amazon Bedrock.

C. Configure AWS Audit Manager as the logs destination for the model.

D. Configure model invocation logging in Amazon EventBridge.

Correct Answer: B

Amazon Bedrock provides an option to enable invocation logging to capture and store the input and output data of the models used. This is essential for monitoring and auditing purposes, particularly when handling customer data. Option B

(Correct): "Enable invocation logging in Amazon Bedrock":This is the correct answer as it directly enables the logging of

all model invocations, ensuring transparency and traceability.

Option A:"Configure AWS CloudTrail" is incorrect because CloudTrail logs API calls but does not provide specific logging for model inputs and outputs. Option C:"Configure AWS Audit Manager" is incorrect as Audit Manager is used for

compliance reporting, not specific invocation logging for AI models. Option D:"Configure model invocation logging in Amazon EventBridge" is incorrect as EventBridge is for event-driven architectures, not specifically designed for logging AI

model inputs and outputs.

AWS AI Practitioner

References:

Amazon Bedrock Logging Capabilities:AWS emphasizes using built-in logging features in Bedrock to maintain data integrity and transparency in model operations.

QUESTION 8

A company has built a chatbot that can respond to natural language questions with images. The company wants to ensure that the chatbot does not return inappropriate or unwanted images.

Which solution will meet these requirements?

A. Implement moderation APIs.

B. Retrain the model with a general public dataset.

C. Perform model validation.

D. Automate user feedback integration.

Correct Answer: A

Moderation APIs, such as Amazon Rekognition\\'s Content Moderation API, can help filter and block inappropriate or unwanted images from being returned by a chatbot. These APIs are specifically designed to detect and manage undesirable

content in images. Option A (Correct): "Implement moderation APIs":This is the correct answer because moderation APIs are designed to identify and filter inappropriate content, ensuring the chatbot does not return unwanted images. Option

B:"Retrain the model with a general public dataset" is incorrect because retraining does not directly prevent inappropriate content from being returned. Option C:"Perform model validation" is incorrect as it ensures model correctness, not

content moderation.

Option D:"Automate user feedback integration" is incorrect because user feedback does not prevent inappropriate images in real-time.

AWS AI Practitioner

References:

AWS Content Moderation Services:AWS provides moderation APIs for filtering unwanted content from applications.

QUESTION 9

An AI practitioner is building a model to generate images of humans in various professions. The AI practitioner discovered that the input data is biased and that specific attributes affect the image generation and create bias in the model. Which technique will solve the problem?

A. Data augmentation for imbalanced classes

B. Model monitoring for class distribution

C. Retrieval Augmented Generation (RAG)

D. Watermark detection for images

Correct Answer: A

Data augmentation for imbalanced classes is the correct technique to address bias in input data affecting image generation.

QUESTION 10

A medical company deployed a disease detection model on Amazon Bedrock. To comply with privacy policies, the company wants to prevent the model from including personal patient information in its responses. The company also wants to receive notification when policy violations occur.

Which solution meets these requirements?

A. Use Amazon Macie to scan the model\\'s output for sensitive data and set up alerts for potential violations.

B. Configure AWS CloudTrail to monitor the model\\'s responses and create alerts for any detected personal information.

C. Use Guardrails for Amazon Bedrock to filter content. Set up Amazon CloudWatch alarms for notification of policy violations.

D. Implement Amazon SageMaker Model Monitor to detect data drift and receive alerts when model quality degrades.

Correct Answer: C

Guardrails for Amazon Bedrock provide mechanisms to filter and control the content generated by models to comply with privacy and policy requirements. Using guardrails ensures that sensitive or personal information is not included in the

model\\'s responses. Additionally, integrating Amazon CloudWatch alarms allows for real-time notification when a policy violation occurs.

Option C (Correct): "Use Guardrails for Amazon Bedrock to filter content. Set up Amazon CloudWatch alarms for notification of policy violations":This is the correct answer because it directly addresses both the prevention of policy violations

and the requirement to receive notifications when such violations occur. Option A:"Use Amazon Macie to scan the model\\'s output for sensitive data" is incorrect because Amazon Macie is designed to monitor data in S3, not to filter real-time

model outputs.

Option B:"Configure AWS CloudTrail to monitor the model\\'s responses" is incorrect because CloudTrail tracks API activity and is not suited for content moderation.

Option D:"Implement Amazon SageMaker Model Monitor to detect data drift" is incorrect because data drift detection does not address content moderation or privacy compliance.

AWS AI Practitioner

References:

Guardrails in Amazon Bedrock:AWS provides guardrails to ensure AI models comply with content policies, and using CloudWatch for alerting integrates monitoring capabilities.

QUESTION 11

A company needs to choose a model from Amazon Bedrock to use internally. The company must identify a model that generates responses in a style that the company\\'s employees prefer.

What should the company do to meet these requirements?

A. Evaluate the models by using built-in prompt datasets.

B. Evaluate the models by using a human workforce and custom prompt datasets.

C. Use public model leaderboards to identify the model.

D. Use the model InvocationLatency runtime metrics in Amazon CloudWatch when trying models.

Correct Answer: B

To determine which model generates responses in a style that the company\\'s employees prefer, the best approach is to use a human workforce to evaluate the models with custom prompt datasets. This method allows for subjective

evaluation based on the specific stylistic preferences of the company\\'s employees, which cannot be effectively assessed through automated methods or pre-built datasets. Option B (Correct): "Evaluate the models by using a human workforce

and custom prompt datasets":This is the correct answer as it directly involves human judgment to evaluate the style and quality of the responses, aligning with employee preferences.

Option A:"Evaluate the models by using built-in prompt datasets" is incorrect because built-in datasets may not capture the company\\'s specific stylistic requirements.

Option C:"Use public model leaderboards to identify the model" is incorrect as leaderboards typically measure model performance on standard benchmarks, not on stylistic preferences.

Option D:"Use the model InvocationLatency runtime metrics in Amazon CloudWatch" is incorrect because latency metrics do not provide any information about the style of the model\\'s responses.

AWS AI Practitioner

References:

Model Evaluation Techniques on AWS:AWS suggests using human evaluators to assess qualitative aspects of model outputs, such as style and tone, to ensure alignment with organizational preferences

---

QUESTION 12

A company is building a chatbot to improve user experience. The company is using a large language model (LLM) from Amazon Bedrock for intent detection. The company wants to use few-shot learning to improve intent detection accuracy.

Which additional data does the company need to meet these requirements?

A. Pairs of chatbot responses and correct user intents

B. Pairs of user messages and correct chatbot responses

C. Pairs of user messages and correct user intents

D. Pairs of user intents and correct chatbot responses

Correct Answer: C

Few-shot learning involves providing a model with a few examples (shots) to learn from. For improving intent detection accuracy in a chatbot using a large language model (LLM), the data should consist of pairs of user messages and their corresponding correct intents.

QUESTION 13

A digital devices company wants to predict customer demand for memory hardware. The company does not have coding experience or knowledge of ML algorithms and needs to develop a data-driven predictive model. The company needs to perform analysis on internal data and external data.

Which solution will meet these requirements?

A. Store the data in Amazon S3. Create ML models and demand forecast predictions by using Amazon SageMaker built-in algorithms that use the data from Amazon S3.

B. Import the data into Amazon SageMaker Data Wrangler. Create ML models and demand forecast predictions by using SageMaker built-in algorithms.

C. Import the data into Amazon SageMaker Data Wrangler. Build ML models and demand forecast predictions by using an Amazon Personalize Trending-Now recipe.

D. Import the data into Amazon SageMaker Canvas. Build ML models and demand forecast predictions by selecting the values in the data from SageMaker Canvas.

Correct Answer: D

Amazon SageMaker Canvas is a visual, no-code machine learning interface that allows users to build machine learning models without having any coding experience or knowledge of machine learning algorithms. It enables users to analyze

internal and external data, and make predictions using a guided interface. Option D (Correct): "Import the data into Amazon SageMaker Canvas. Build ML models and demand forecast predictions by selecting the values in the data from

SageMaker Canvas":This is the correct answer because SageMaker Canvas is designed for users without coding experience, providing a visual interface to build predictive models with ease.

Option A:"Store the data in Amazon S3 and use SageMaker built-in algorithms" is incorrect because it requires coding knowledge to interact with SageMaker\\'s built- in algorithms.

Option B:"Import the data into Amazon SageMaker Data Wrangler" is incorrect. Data Wrangler is primarily for data preparation and not directly focused on creating ML models without coding.

Option C:"Use Amazon Personalize Trending-Now recipe" is incorrect as Amazon Personalize is for building recommendation systems, not for general demand forecasting.

AWS AI Practitioner

References:

Amazon SageMaker Canvas Overview:AWS documentation emphasizes Canvas as a no-code solution for building machine learning models, suitable for business analysts and users with no coding experience.

QUESTION 14

Which feature of Amazon OpenSearch Service gives companies the ability to build vector database applications?

A. Integration with Amazon S3 for object storage

B. Support for geospatial indexing and queries

C. Scalable index management and nearest neighbor search capability

D. Ability to perform real-time analysis on streaming data

Correct Answer: C

Amazon OpenSearch Service (formerly Amazon Elasticsearch Service) has introduced capabilities to support vector search, which allows companies to build vector database applications. This is particularly useful in machine learning, where

vector representations (embeddings) of data are often used to capture semantic meaning. Scalable index management and nearest neighbor search capabilityare the core features enabling vector database functionalities in OpenSearch. The

service allows users to index high-dimensional vectors and perform efficient nearest neighbor searches, which are crucial for tasks such as recommendation systems, anomaly detection, and semantic search.

Here is why option C is the correct answer:

Scalable Index Management:OpenSearch Service supports scalable indexing of vector data. This means you can index a large volume of high-dimensional vectors and manage these indexes in a cost-effective and performance-optimized

way. The service leverages underlying AWS infrastructure to ensure that indexing scales seamlessly with data size.

Nearest Neighbor Search Capability:OpenSearch Service\\'s nearest neighbor search capability allows for fast and efficient searches over vector data. This is essential for applicationslike product recommendation engines, where the system

needs to quickly find the most similar items based on a user\\'s query or behavior.

AWS AI Practitioner References:

The other options do not directly relate to building vector database applications:

A. Integration with Amazon S3 for object storageis about storing data objects, not vector-based searching or indexing.

B. Support for geospatial indexing and queriesis related to location-based data, not vectors used in machine learning.

D. Ability to perform real-time analysis on streaming datarelates to analyzing incoming data streams, which is different from the vector search capabilities.

QUESTION 15

A company is building a contact center application and wants to gain insights from customer conversations. The company wants to analyze and extract key information from the audio of the customer calls. Which solution meets these requirements?

A. Build a conversational chatbot by using Amazon Lex.

B. Transcribe call recordings by using Amazon Transcribe.

C. Extract information from call recordings by using Amazon SageMaker Model Monitor.

D. Create classification labels by using Amazon Comprehend.

Correct Answer: B

Amazon Transcribe is the correct solution for converting audio from customer calls into text, allowing the company to analyze and extract key information from the conversations.