Amazon-Web-Services

Exam Questions AIF-C01

AWS Certified AI Practitioner

NEW QUESTION 1
A company wants to use a large language model (LLM) on Amazon Bedrock for sentiment analysis. The company wants to know how much information can fit into one prompt.
Which consideration will inform the company's decision?

A. Temperature
B. Context window
C. Batch size
D. Model size

Answer: B

Explanation:
The context window determines how much information can fit into a single prompt when using a large language model (LLM) like those on Amazon Bedrock.
? Context Window:
? Why Option B is Correct:
? Why Other Options are Incorrect:

NEW QUESTION 2
A company wants to display the total sales for its top-selling products across various retail locations in the past 12 months.
Which AWS solution should the company use to automate the generation of graphs?

A. Amazon Q in Amazon EC2
B. Amazon Q Developer
C. Amazon Q in Amazon QuickSight
D. Amazon Q in AWS Chatbot

Answer: C

Explanation:
Amazon QuickSight is a fully managed business intelligence (BI) service that allows users to create and publish interactive dashboards that include visualizations like graphs, charts, and tables. "Amazon Q" is the natural language query feature within Amazon QuickSight. It enables users to ask questions about their data in natural language and receive visual responses such as graphs.
? Option C (Correct): "Amazon Q in Amazon QuickSight": This is the correct answer
because Amazon QuickSight Q is specifically designed to allow users to explore their data through natural language queries, and it can automatically generate graphs to display sales data and other metrics. This makes it an ideal choice for the company to automate the generation of graphs showing total sales for its top-selling products across various retail locations.
? Option A, B, and D: These options are incorrect:
AWS AI Practitioner References:
? Amazon QuickSight Q is designed to provide insights from data by using natural language queries, making it a powerful tool for generating automated graphs and visualizations directly from queried data.
? Business Intelligence (BI) on AWS: AWS services such as Amazon QuickSight
provide business intelligence capabilities, including automated reporting and visualization features, which are ideal for companies seeking to visualize data like sales trends over time.

NEW QUESTION 3
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)

Answer: A

Explanation:
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.
? Confusion Matrix:
? Why Option A is Correct:
? Why Other Options are Incorrect:

NEW QUESTION 4
A company has documents that are missing some words because of a database error. The company wants to build an ML model that can suggest potential words to fill in the missing text.
Which type of model meets this requirement?

A. Topic modeling
B. Clustering models
C. Prescriptive ML models
D. BERT-based models

Answer: D

Explanation:
BERT-based models (Bidirectional Encoder Representations from Transformers) are suitable for tasks that involve understanding the context of words in a sentence and suggesting missing words. These models use bidirectional training, which considers the context from both directions (left and right of the missing

word) to predict the appropriate word to fill in the gaps.
? BERT-based Models:
? Why Option D is Correct:
? Why Other Options are Incorrect:

NEW QUESTION 5
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 conten
D. Set up Amazon CloudWatch alarms for notification of policy violations.
E. Implement Amazon SageMaker Model Monitor to detect data drift and receive alerts when model quality degrades.

Answer: C

Explanation:
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.

NEW QUESTION 6
A company has a foundation model (FM) that was customized by using Amazon Bedrock to answer customer queries about products. The company wants to validate the model's responses to new types of queries. The company needs to upload a new dataset that Amazon Bedrock can use for validation.
Which AWS service meets these requirements?

A. Amazon S3
B. Amazon Elastic Block Store (Amazon EBS)
C. Amazon Elastic File System (Amazon EFS)
D. AWS Snowcone

Answer: A

Explanation:
Amazon S3 is the optimal choice for storing and uploading datasets used for machine learning model validation and training. It offers scalable, durable, and secure storage, making it ideal for holding datasets required by Amazon Bedrock for validation purposes.
? Option A (Correct): "Amazon S3": This is the correct answer because Amazon S3
is widely used for storing large datasets that are accessed by machine learning models, including those in Amazon Bedrock.
? Option B: "Amazon Elastic Block Store (Amazon EBS)" is incorrect because EBS
is a block storage service for use with Amazon EC2, not for directly storing datasets for Amazon Bedrock.
? Option C: "Amazon Elastic File System (Amazon EFS)" is incorrect as it is
primarily used for file storage with shared access by multiple instances.
? Option D: "AWS Snowcone" is incorrect because it is a physical device for offline data transfer, not suitable for directly providing data to Amazon Bedrock.
AWS AI Practitioner References:
? Storing and Managing Datasets on AWS for Machine Learning: AWS recommends using S3 for storing and managing datasets required for ML model training and validation.

NEW QUESTION 7
Which strategy evaluates the accuracy of a foundation model (FM) that is used in image classification tasks?

A. Calculate the total cost of resources used by the model.
B. Measure the model's accuracy against a predefined benchmark dataset.
C. Count the number of layers in the neural network.
D. Assess the color accuracy of images processed by the model.

Answer: B

Explanation:
Measuring the model's accuracy against a predefined benchmark dataset is the correct strategy to evaluate the accuracy of a foundation model (FM) used in image classification tasks.
? Model Accuracy Evaluation:
? Why Option B is Correct:
? Why Other Options are Incorrect:

NEW QUESTION 8

What does an F1 score measure in the context of foundation model (FM) performance?

A. Model precision and recall.
B. Model speed in generating responses.
C. Financial cost of operating the model.
D. Energy efficiency of the model's computations.

Answer: A

Explanation:
The F1 score is the harmonic mean of precision and recall, making it a balanced metric for evaluating model performance when there is an imbalance between false positives and false negatives. Speed, cost, and energy efficiency are unrelated to the F1 score. References: AWS Foundation Models Guide.

NEW QUESTION 9
A company built an AI-powered resume screening system. The company used a large dataset to train the model. The dataset contained resumes that were not representative of all demographics. Which core dimension of responsible AI does this scenario present?

A. Fairness.
B. Explainability.
C. Privacy and security.
D. Transparency.

Answer: A

Explanation:
Fairness refers to the absence of bias in AI models. Using non- representative datasets leads to biased predictions, affecting specific demographics unfairly. Explainability, privacy, and transparency are important but not directly related to this scenario. References: AWS Responsible AI Framework.

NEW QUESTION 10
A company wants to classify human genes into 20 categories based on gene characteristics. The company needs an ML algorithm to document how the inner mechanism of the model affects the output.
Which ML algorithm meets these requirements?

A. Decision trees
B. Linear regression
C. Logistic regression
D. Neural networks

Answer: A

Explanation:
Decision trees are an interpretable machine learning algorithm that clearly documents the decision-making process by showing how each input feature affects the output. This transparency is particularly useful when explaining how the model arrives at a certain decision, making it suitable for classifying genes into categories.
? Option A (Correct): "Decision trees": This is the correct answer because decision
trees provide a clear and interpretable representation of how input features influence the model's output, making it ideal for understanding the inner mechanisms affecting predictions.
? Option B: "Linear regression" is incorrect because it is used for regression tasks,
not classification.
? Option C: "Logistic regression" is incorrect as it does not provide the same level of interpretability in documenting decision-making processes.
? Option D: "Neural networks" is incorrect because they are often considered "black boxes" and do not easily explain how they arrive at their outputs.
AWS AI Practitioner References:
? Interpretable Machine Learning Models on AWS: AWS supports using interpretable models, such as decision trees, for tasks that require clear documentation of how input data affects output decisions.

NEW QUESTION 10
What does an F1 score measure in the context of foundation model (FM) performance?

A. Model precision and recall
B. Model speed in generating responses
C. Financial cost of operating the model
D. Energy efficiency of the model's computations

Answer: A

Explanation:
The F1 score is a metric used to evaluate the performance of a classification model by considering both precision and recall. Precision measures the accuracy of positive predictions (i.e., the proportion of true positive predictions among all positive predictions made by the model), while recall measures the model's ability to identify all relevant positive instances (i.e., the proportion of true positive predictions among all actual positive instances). The F1 score is the harmonic mean of precision and recall, providing a single metric that balances both concerns. This is particularly useful when dealing with imbalanced datasets or when the cost of false positives and false negatives is significant. Options B, C, and D pertain to other aspects of model performance but are not related to the F1 score.
Reference: AWS Certified AI Practitioner Exam Guide

NEW QUESTION 11
An education provider is building a question and answer application that uses a generative AI model to explain complex concepts. The education provider wants to automatically change the style of the model response depending on who is asking the question. The education provider will give the model the age range of the user who has asked the question.
Which solution meets these requirements with the LEAST implementation effort?

A. Fine-tune the model by using additional training data that is representative of the various age ranges that the application will support.

B. Add a role description to the prompt context that instructs the model of the age range that the response should target.
C. Use chain-of-thought reasoning to deduce the correct style and complexity for a response suitable for that user.
D. Summarize the response text depending on the age of the user so that younger users receive shorter responses.

Answer: B

Explanation:
Adding a role description to the prompt context is a straightforward way to instruct the generative AI model to adjust its response style based on the user's age range. This method requires minimal implementation effort as it does not involve additional training or complex logic.
? Option B (Correct): "Add a role description to the prompt context that instructs the model of the age range that the response should target": This is the correct answer because it involves the least implementation effort while effectively guiding the
model to tailor responses according to the age range.
? Option A: "Fine-tune the model by using additional training data" is incorrect because it requires significant effort in gathering data and retraining the model.
? Option C: "Use chain-of-thought reasoning" is incorrect as it involves complex reasoning that may not directly address the need to adjust response style based on age.
? Option D: "Summarize the response text depending on the age of the user" is incorrect because it involves additional processing steps after generating the initial response, increasing complexity.
AWS AI Practitioner References:
? Prompt Engineering Techniques on AWS: AWS recommends using prompt context effectively to guide generative models in providing tailored responses based on specific user attributes.

NEW QUESTION 13
A company is using an Amazon Bedrock base model to summarize documents for an internal use case. The company trained a custom model to improve the summarization quality.
Which action must the company take to use the custom model through Amazon Bedrock?

A. Purchase Provisioned Throughput for the custom model.
B. Deploy the custom model in an Amazon SageMaker endpoint for real-time inference.
C. Register the model with the Amazon SageMaker Model Registry.
D. Grant access to the custom model in Amazon Bedrock.

Answer: B

Explanation:
To use a custom model that has been trained to improve summarization quality, the company must deploy the model on an Amazon SageMaker endpoint. This allows the model to be used for real-time inference through Amazon Bedrock or other AWS services. By deploying the model in SageMaker, the custom model can be accessed programmatically via API calls, enabling integration with Amazon Bedrock.
? Option B (Correct): "Deploy the custom model in an Amazon SageMaker endpoint
for real-time inference": This is the correct answer because deploying the model on SageMaker enables it to serve real-time predictions and be integrated with Amazon Bedrock.
? Option A: "Purchase Provisioned Throughput for the custom model" is incorrect
because provisioned throughput is related to database or storage services, not model deployment.
? Option C: "Register the model with the Amazon SageMaker Model Registry" is
incorrect because while the model registry helps with model management, it does not make the model accessible for real-time inference.
? Option D: "Grant access to the custom model in Amazon Bedrock" is incorrect
because Bedrock does not directly manage custom model access; it relies on deployed endpoints like those in SageMaker.
AWS AI Practitioner References:
? Amazon SageMaker Endpoints: AWS recommends deploying models to SageMaker endpoints to use them for real-time inference in various applications.

NEW QUESTION 15
A company built a deep learning model for object detection and deployed the model to production.
Which AI process occurs when the model analyzes a new image to identify objects?

A. Training
B. Inference
C. Model deployment
D. Bias correction

Answer: B

Explanation:
Inference is the correct answer because it is the AI process that occurs when a deployed model analyzes new data (such as an image) to make predictions or identify objects.
? Inference:
? Why Option B is Correct:
? Why Other Options are Incorrect:

NEW QUESTION 20
A company wants to use language models to create an application for inference on edge devices. The inference must have the lowest latency possible.
Which solution will meet these requirements?

A. Deploy optimized small language models (SLMs) on edge devices.
B. Deploy optimized large language models (LLMs) on edge devices.
C. Incorporate a centralized small language model (SLM) API for asynchronous communication with edge devices.
D. Incorporate a centralized large language model (LLM) API for asynchronous communication with edge devices.

Answer: A

Explanation:
To achieve the lowest latency possible for inference on edge devices, deploying optimized small language models (SLMs) is the most effective solution. SLMs

require fewer
resources and have faster inference times, making them ideal for deployment on edge devices where processing power and memory are limited.
? Option A (Correct): "Deploy optimized small language models (SLMs) on edge
devices": This is the correct answer because SLMs provide fast inference with low latency, which is crucial for edge deployments.
? Option B: "Deploy optimized large language models (LLMs) on edge devices" is
incorrect because LLMs are resource-intensive and may not perform well on edge devices due to their size and computational demands.
? Option C: "Incorporate a centralized small language model (SLM) API for
asynchronous communication with edge devices" is incorrect because it introduces network latency due to the need for communication with a centralized server.
? Option D: "Incorporate a centralized large language model (LLM) API for
asynchronous communication with edge devices" is incorrect for the same reason, with even greater latency due to the larger model size.
AWS AI Practitioner References:
? Optimizing AI Models for Edge Devices on AWS: AWS recommends using small, optimized models for edge deployments to ensure minimal latency and efficient performance.

NEW QUESTION 23
A retail store wants to predict the demand for a specific product for the next few weeks by using the Amazon SageMaker DeepAR forecasting algorithm.
Which type of data will meet this requirement?

A. Text data
B. Image data
C. Time series data
D. Binary data

Answer: C

Explanation:
Amazon SageMaker's DeepAR is a supervised learning algorithm designed for forecasting scalar (one-dimensional) time series data. Time series data consists of sequences of data points indexed in time order, typically with consistent intervals between them. In the context of a retail store aiming to predict product demand, relevant time series data might include historical sales figures, inventory levels, or related metrics recorded over regular time intervals (e.g., daily or weekly). By training the DeepAR model on this historical time series data, the store can generate forecasts for future product demand. This capability is
particularly useful for inventory management, staffing, and supply chain optimization. Other data types, such as text, image, or binary data, are not suitable for time series forecasting tasks and would not be appropriate inputs for the DeepAR algorithm.
Reference: Amazon SageMaker DeepAR Algorithm

NEW QUESTION 28
An AI company periodically evaluates its systems and processes with the help of independent software vendors (ISVs). The company needs to receive email message notifications when an ISV's compliance reports become available.
Which AWS service can the company use to meet this requirement?

A. AWS Audit Manager
B. AWS Artifact
C. AWS Trusted Advisor
D. AWS Data Exchange

Answer: D

Explanation:
AWS Data Exchange is a service that allows companies to securely exchange data with third parties, such as independent software vendors (ISVs). AWS Data Exchange can be configured to provide notifications, including email notifications, when new datasets or compliance reports become available.
? Option D (Correct): "AWS Data Exchange": This is the correct answer because it
enables the company to receive notifications, including email messages, when ISVs' compliance reports are available.
? Option A: "AWS Audit Manager" is incorrect because it focuses on assessing an
organization's own compliance, not receiving third-party compliance reports.
? Option B: "AWS Artifact" is incorrect as it provides access to AWS??s compliance reports, not ISVs'.
? Option C: "AWS Trusted Advisor" is incorrect as it offers optimization and best practices guidance, not compliance report notifications.
AWS AI Practitioner References:
? AWS Data Exchange Documentation: AWS explains how Data Exchange allows organizations to subscribe to third-party data and receive notifications when updates are available.

NEW QUESTION 30
A loan company is building a generative AI-based solution to offer new applicants discounts based on specific business criteria. The company wants to build and use an AI model responsibly to minimize bias that could negatively affect some customers.
Which actions should the company take to meet these requirements? (Select TWO.)

A. Detect imbalances or disparities in the data.
B. Ensure that the model runs frequently.
C. Evaluate the model's behavior so that the company can provide transparency to stakeholders.
D. Use the Recall-Oriented Understudy for Gisting Evaluation (ROUGE) technique to ensure that the model is 100% accurate.
E. Ensure that the model's inference time is within the accepted limits.

Answer: AC

Explanation:
To build and use an AI model responsibly, especially in sensitive applications like loan approvals, it's crucial to address potential biases and ensure transparency:
? Detect imbalances or disparities in the data (Option A): Analyzing the training data
for imbalances or disparities is essential. Imbalanced data can lead to models that are biased towards the majority class, potentially disadvantaging certain groups. By identifying and mitigating these imbalances, the company can reduce the risk of biased predictions.
? Evaluate the model's behavior to provide transparency to stakeholders (Option C):
Regularly assessing the model's outputs and decision-making processes allows the company to understand how decisions are made. This evaluation fosters transparency, enabling the company to explain model behavior to stakeholders

and ensure that the model operates as intended without unintended biases. Options B, D, and E, while relevant to model performance and evaluation, do not directly address the responsible use of AI concerning bias and transparency.
Reference: AWS Certified AI Practitioner Exam Guide

NEW QUESTION 33
Which metric measures the runtime efficiency of operating AI models?

A. Customer satisfaction score (CSAT)
B. Training time for each epoch
C. Average response time
D. Number of training instances

Answer: C

Explanation:
The average response time is the correct metric for measuring the runtime efficiency of operating AI models.
? Average Response Time:
? Why Option C is Correct:
? Why Other Options are Incorrect:

NEW QUESTION 37
An accounting firm wants to implement a large language model (LLM) to automate document processing. The firm must proceed responsibly to avoid potential harms.
What should the firm do when developing and deploying the LLM? (Select TWO.)

A. Include fairness metrics for model evaluation.
B. Adjust the temperature parameter of the model.
C. Modify the training data to mitigate bias.
D. Avoid overfitting on the training data.
E. Apply prompt engineering techniques.

Answer: AC

Explanation:
To implement a large language model (LLM) responsibly, the firm should focus on fairness and mitigating bias, which are critical for ethical AI deployment.
? A. Include Fairness Metrics for Model Evaluation:
? C. Modify the Training Data to Mitigate Bias:
? Why Other Options are Incorrect:

NEW QUESTION 39
A company is building a customer service chatbot. The company wants the chatbot to improve its responses by learning from past interactions and online resources.
Which AI learning strategy provides this self-improvement capability?

A. Supervised learning with a manually curated dataset of good responses and bad responses
B. Reinforcement learning with rewards for positive customer feedback
C. Unsupervised learning to find clusters of similar customer inquiries
D. Supervised learning with a continuously updated FAQ database

Answer: B

Explanation:
Reinforcement learning allows a model to learn and improve over time based on feedback from its environment. In this case, the chatbot can improve its responses by being rewarded for positive customer feedback, which aligns well with the goal of self- improvement based on past interactions and new information.
? Option B (Correct): "Reinforcement learning with rewards for positive customer
feedback": This is the correct answer as reinforcement learning enables the chatbot to learn from feedback and adapt its behavior accordingly, providing self-improvement capabilities.
? Option A: "Supervised learning with a manually curated dataset" is incorrect
because it does not support continuous learning from new interactions.
? Option C: "Unsupervised learning to find clusters of similar customer inquiries" is incorrect because unsupervised learning does not provide a mechanism for improving responses based on feedback.
? Option D: "Supervised learning with a continuously updated FAQ database" is incorrect because it still relies on manually curated data rather than self-improvement from feedback.
AWS AI Practitioner References:
? Reinforcement Learning on AWS: AWS provides reinforcement learning
frameworks that can be used to train models to improve their performance based on feedback.

NEW QUESTION 41
A company wants to create an application by using Amazon Bedrock. The company has a limited budget and prefers flexibility without long-term commitment.
Which Amazon Bedrock pricing model meets these requirements?

A. On-Demand
B. Model customization
C. Provisioned Throughput
D. Spot Instance

Answer: A

Explanation:

Amazon Bedrock offers an on-demand pricing model that provides flexibility without long- term commitments. This model allows companies to pay only for the resources they use, which is ideal for a limited budget and offers flexibility.
? Option A (Correct): "On-Demand": This is the correct answer because on-demand
pricing allows the company to use Amazon Bedrock without any long-term commitments and to manage costs according to their budget.
? Option B: "Model customization" is a feature, not a pricing model.
? Option C: "Provisioned Throughput" involves reserving capacity ahead of time, which might not offer the desired flexibility and could lead to higher costs if the capacity is not fully used.
? Option D: "Spot Instance" is a pricing model for EC2 instances and does not apply to Amazon Bedrock.
AWS AI Practitioner References:
? AWS Pricing Models for Flexibility: On-demand pricing is a key AWS model for services that require flexibility and no long-term commitment, ensuring cost-effectiveness for projects with variable usage patterns.

NEW QUESTION 43
A social media company wants to use a large language model (LLM) for content moderation. The company wants to evaluate the LLM outputs for bias and potential discrimination against specific groups or individuals.
Which data source should the company use to evaluate the LLM outputs with the LEAST administrative effort?

A. User-generated content
B. Moderation logs
C. Content moderation guidelines
D. Benchmark datasets

Answer: D

Explanation:
Benchmark datasets are pre-validated datasets specifically designed to evaluate machine learning models for bias, fairness, and potential discrimination. These datasets are the most efficient tool for assessing an LLM??s performance against known standards with
minimal administrative effort.
? Option D (Correct): "Benchmark datasets": This is the correct answer because using standardized benchmark datasets allows the company to evaluate model outputs for bias with minimal administrative overhead.
? Option A: "User-generated content" is incorrect because it is unstructured and would require significant effort to analyze for bias.
? Option B: "Moderation logs" is incorrect because they represent historical data and do not provide a standardized basis for evaluating bias.
? Option C: "Content moderation guidelines" is incorrect because they provide qualitative criteria rather than a quantitative basis for evaluation.
AWS AI Practitioner References:
? Evaluating AI Models for Bias on AWS: AWS supports using benchmark datasets to assess model fairness and detect potential bias efficiently.

NEW QUESTION 48
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

Answer: B

Explanation:
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.
? AWS CloudTrail:
? Why Option B is Correct:
? Why Other Options are Incorrect:
Thus, B is the correct answer for identifying unauthorized users attempting to access Amazon Bedrock.

NEW QUESTION 51
A company wants to use generative AI to increase developer productivity and software development. The company wants to use Amazon Q Developer.
What can Amazon Q Developer do to help the company meet these requirements?

A. Create software snippets, reference tracking, and open-source license tracking.
B. Run an application without provisioning or managing servers.
C. Enable voice commands for coding and providing natural language search.
D. Convert audio files to text documents by using ML models.

Answer: A

Explanation:
Amazon Q Developer is a tool designed to assist developers in increasing productivity by generating code snippets, managing reference tracking, and handling open-source license tracking. These features help developers by automating parts of the software development process.
? Option A (Correct): "Create software snippets, reference tracking, and open-
source license tracking": This is the correct answer because these are key features that help developers streamline and automate tasks, thus improving productivity.
? Option B: "Run an application without provisioning or managing servers" is
incorrect as it refers to AWS Lambda or AWS Fargate, not Amazon Q Developer.
? Option C: "Enable voice commands for coding and providing natural language search" is incorrect because this is not a function of Amazon Q Developer.
? Option D: "Convert audio files to text documents by using ML models" is incorrect as this refers to Amazon Transcribe, not Amazon Q Developer.
AWS AI Practitioner References:

? Amazon Q Developer Features: AWS documentation outlines how Amazon Q Developer supports developers by offering features that reduce manual effort and improve efficiency.

NEW QUESTION 56
An AI practitioner trained a custom model on Amazon Bedrock by using a training dataset that contains confidential data. The AI practitioner wants to ensure that the custom model does not generate inference responses based on confidential data.
How should the AI practitioner prevent responses based on confidential data?

A. Delete the custom mode
B. Remove the confidential data from the training datase
C. Retrain the custom model.
D. Mask the confidential data in the inference responses by using dynamic data masking.
E. Encrypt the confidential data in the inference responses by using Amazon SageMaker.
F. Encrypt the confidential data in the custom model by using AWS Key Management Service (AWS KMS).

Answer: A

Explanation:
When a model is trained on a dataset containing confidential or sensitive data, the model may inadvertently learn patterns from this data, which could then be reflected in its inference responses. To ensure that a model does not generate responses based on confidential data, the most effective approach is to remove the confidential data from the training dataset and then retrain the model.
Explanation of Each Option:
? Option A (Correct): "Delete the custom model. Remove the confidential data from the training dataset. Retrain the custom model."This option is correct because it directly addresses the core issue: the model has been trained on confidential data. The only way to ensure that the model does not produce inferences based on this data is to remove the confidential information from the training dataset and then retrain the model from scratch. Simply deleting the model and retraining it ensures that no confidential data is learned or retained by the model. This approach follows the best practices recommended by AWS for handling sensitive data when using machine learning services like Amazon Bedrock.
? Option B: "Mask the confidential data in the inference responses by using dynamic data masking."This option is incorrect because dynamic data masking is typically used to mask or obfuscate sensitive data in a database. It does not address the core problem of the model being trained on confidential data. Masking data in inference responses does not prevent the model from using confidential data it learned during training.
? Option C: "Encrypt the confidential data in the inference responses by using Amazon SageMaker."This option is incorrect because encrypting the inference responses does not prevent the model from generating outputs based on confidential data. Encryption only secures the data at rest or in transit but does not affect the model's underlying knowledge or training process.
? Option D: "Encrypt the confidential data in the custom model by using AWS Key Management Service (AWS KMS)."This option is incorrect as well because encrypting the data within the model does not prevent the model from generating responses based on the confidential data it learned during training. AWS KMS can encrypt data, but it does not modify the learning that the model has already performed.
AWS AI Practitioner References:
? Data Handling Best Practices in AWS Machine Learning: AWS advises practitioners to carefully handle training data, especially when it involves sensitive or confidential information. This includes preprocessing steps like data anonymization or removal of sensitive data before using it to train machine learning models.
? Amazon Bedrock and Model Training Security: Amazon Bedrock provides foundational models and customization capabilities, but any training involving sensitive data should follow best practices, such as removing or anonymizing confidential data to prevent unintended data leakage.

NEW QUESTION 58
A medical company is customizing a foundation model (FM) for diagnostic purposes. The company needs the model to be transparent and explainable to meet regulatory requirements.
Which solution will meet these requirements?

A. Configure the security and compliance by using Amazon Inspector.
B. Generate simple metrics, reports, and examples by using Amazon SageMaker Clarify.
C. Encrypt and secure training data by using Amazon Macie.
D. Gather more dat
E. Use Amazon Rekognition to add custom labels to the data.

Answer: B

Explanation:
Amazon SageMaker Clarify provides transparency and explainability for machine learning models by generating metrics, reports, and examples that help to understand model predictions. For a medical company that needs a foundation model to be transparent and explainable to meet regulatory requirements, SageMaker Clarify is the most suitable solution.
? Amazon SageMaker Clarify:
? Why Option B is Correct:
? Why Other Options are Incorrect:
Thus, B is the correct answer for meeting transparency and explainability requirements for the foundation model

NEW QUESTION 60
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

Answer: B

Explanation:
Sampling bias is the correct type of bias affecting the model output when it disproportionately flags people from a specific ethnic group.
? Sampling Bias:
? Why Option B is Correct:

? Why Other Options are Incorrect:

NEW QUESTION 64
A company is implementing the Amazon Titan foundation model (FM) by using Amazon Bedrock. The company needs to supplement the model by using relevant data from the company's private data sources.
Which solution will meet this requirement?

A. Use a different FM
B. Choose a lower temperature value
C. Create an Amazon Bedrock knowledge base
D. Enable model invocation logging

Answer: C

Explanation:
Creating an Amazon Bedrock knowledge base allows the integration of external or private data sources with a foundation model (FM) like Amazon Titan. This integration helps supplement the model with relevant data from the company's private data sources to enhance its responses.
? Option C (Correct): "Create an Amazon Bedrock knowledge base": This is the
correct answer as it enables the company to incorporate private data into the FM to improve its effectiveness.
? Option A: "Use a different FM" is incorrect because it does not address the need to
supplement the current model with private data.
? Option B: "Choose a lower temperature value" is incorrect as it affects output randomness, not the integration of private data.
? Option D: "Enable model invocation logging" is incorrect because logging does not help in supplementing the model with additional data.
AWS AI Practitioner References:
? Amazon Bedrock and Knowledge Integration: AWS explains how creating a knowledge base allows Amazon Bedrock to use external data sources to improve the FM??s relevance and accuracy.

NEW QUESTION 69
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 Wrangle
C. Create ML models and demand forecast predictions by using SageMaker built-in algorithms.
D. Import the data into Amazon SageMaker Data Wrangle
E. Build ML models and demand forecast predictions by using an Amazon Personalize Trending-Now recipe.
F. Import the data into Amazon SageMaker Canva
G. Build ML models and demand forecast predictions by selecting the values in the data from SageMaker Canvas.

Answer: D

Explanation:
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.

NEW QUESTION 70
An AI practitioner is using a large language model (LLM) to create content for marketing campaigns. The generated content sounds plausible and factual but is incorrect.
Which problem is the LLM having?

A. Data leakage
B. Hallucination
C. Overfitting
D. Underfitting

Answer: B

Explanation:
In the context of AI, "hallucination" refers to the phenomenon where a model generates outputs that are plausible-sounding but are not grounded in reality or the training data. This problem often occurs with large language models (LLMs) when they create information that sounds correct but is actually incorrect or fabricated.
? Option B (Correct): "Hallucination": This is the correct answer because the
problem described involves generating content that sounds factual but is incorrect, which is characteristic of hallucination in generative AI models.
? Option A: "Data leakage" is incorrect as it involves the model accidentally learning
from data it shouldn't have access to, which does not match the problem of generating incorrect content.
? Option C: "Overfitting" is incorrect because overfitting refers to a model that has
learned the training data too well, including noise, and performs poorly on new data.

? Option D: "Underfitting" is incorrect because underfitting occurs when a model is
too simple to capture the underlying patterns in the data, which is not the issue here.
AWS AI Practitioner References:
? Large Language Models on AWS: AWS discusses the challenge of hallucination in large language models and emphasizes techniques to mitigate it, such as using guardrails and fine-tuning.

NEW QUESTION 75
A company is building an ML model to analyze archived data. The company must perform inference on large datasets that are multiple GBs in size. The company does not need to access the model predictions immediately.
Which Amazon SageMaker inference option will meet these requirements?

A. Batch transform
B. Real-time inference
C. Serverless inference
D. Asynchronous inference

Answer: A

Explanation:
Batch transform in Amazon SageMaker is designed for offline processing of large datasets. It is ideal for scenarios where immediate predictions are not required, and the inference can be done on large datasets that are multiple gigabytes in size. This method processes data in batches, making it suitable for analyzing archived data without the need for real- time access to predictions.
? Option A (Correct): "Batch transform": This is the correct answer because batch
transform is optimized for handling large datasets and is suitable when immediate access to predictions is not required.
? Option B: "Real-time inference" is incorrect because it is used for low-latency, real-
time prediction needs, which is not required in this case.
? Option C: "Serverless inference" is incorrect because it is designed for small-scale, intermittent inference requests, not for large batch processing.
? Option D: "Asynchronous inference" is incorrect because it is used when immediate predictions are required, but with high throughput, whereas batch transform is more suitable for very large datasets.
AWS AI Practitioner References:
? Batch Transform on AWS SageMaker: AWS recommends using batch transform for large datasets when real-time processing is not needed, ensuring cost-effectiveness and scalability.

NEW QUESTION 79
An e-commerce company wants to build a solution to determine customer sentiments based on written customer reviews of products.
Which AWS services meet these requirements? (Select TWO.)

A. Amazon Lex
B. Amazon Comprehend
C. Amazon Polly
D. Amazon Bedrock
E. Amazon Rekognition

Answer: BD

Explanation:
To determine customer sentiments based on written customer reviews, the company can use Amazon Comprehend and Amazon Bedrock.
? Amazon Comprehend:
? Amazon Bedrock:
? Why Other Options are Incorrect:

NEW QUESTION 82
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.

Answer: B

Explanation:
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.
? Amazon Transcribe:
? Why Option B is Correct:
? Why Other Options are Incorrect:

NEW QUESTION 85
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

Answer: C

Explanation:
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.
? Few-shot Learning for Intent Detection:
? Why Option C is Correct:
? Why Other Options are Incorrect:

NEW QUESTION 87
A company has built a solution by using generative AI. The solution uses large language models (LLMs) to translate training manuals from English into other languages. The company wants to evaluate the accuracy of the solution by examining the text generated for the manuals.
Which model evaluation strategy meets these requirements?

A. Bilingual Evaluation Understudy (BLEU)
B. Root mean squared error (RMSE)
C. Recall-Oriented Understudy for Gisting Evaluation (ROUGE)
D. F1 score

Answer: A

Explanation:
BLEU (Bilingual Evaluation Understudy) is a metric used to evaluate the accuracy of machine-generated translations by comparing them against reference translations. It is commonly used for translation tasks to measure how close the generated output is to professional human translations.
? Option A (Correct): "Bilingual Evaluation Understudy (BLEU)": This is the correct answer because BLEU is specifically designed to evaluate the quality of translations, making it suitable for the company's use case.
? Option B: "Root mean squared error (RMSE)" is incorrect because RMSE is used for regression tasks to measure prediction errors, not translation quality.
? Option C: "Recall-Oriented Understudy for Gisting Evaluation (ROUGE)" is incorrect as it is used to evaluate text summarization, not translation.
? Option D: "F1 score" is incorrect because it is typically used for classification tasks, not for evaluating translation accuracy.
AWS AI Practitioner References:
? Model Evaluation Metrics on AWS: AWS supports various metrics like BLEU for specific use cases, such as evaluating machine translation models.

NEW QUESTION 89
A company needs to train an ML model to classify images of different types of animals. The company has a large dataset of labeled images and will not label more data. Which type of learning should the company use to train the model?

A. Supervised learning.
B. Unsupervised learning.
C. Reinforcement learning.
D. Active learning.

Answer: A

Explanation:
Supervised learning is appropriate when the dataset is labeled. The model uses this data to learn patterns and classify images. Unsupervised learning, reinforcement learning, and active learning are not suitable since they either require unlabeled data or different problem settings. References: AWS Machine Learning Best Practices.

NEW QUESTION 92
A pharmaceutical company wants to analyze user reviews of new medications and provide a concise overview for each medication. Which solution meets these requirements?

A. Create a time-series forecasting model to analyze the medication reviews by using Amazon Personalize.
B. Create medication review summaries by using Amazon Bedrock large language models (LLMs).
C. Create a classification model that categorizes medications into different groups by using Amazon SageMaker.
D. Create medication review summaries by using Amazon Rekognition.

Answer: B

Explanation:
Amazon Bedrock provides large language models (LLMs) that are optimized for natural language understanding and text summarization tasks, making it the best choice for creating concise summaries of user reviews. Time-series forecasting, classification, and image analysis (Rekognition) are not suitable for summarizing textual data. References: AWS Bedrock Documentation.

NEW QUESTION 96
A company is using few-shot prompting on a base model that is hosted on Amazon Bedrock. The model currently uses 10 examples in the prompt. The model is invoked once daily and is performing well. The company wants to lower the monthly cost.
Which solution will meet these requirements?

A. Customize the model by using fine-tuning.
B. Decrease the number of tokens in the prompt.
C. Increase the number of tokens in the prompt.
D. Use Provisioned Throughput.

Answer: B

Explanation:
Decreasing the number of tokens in the prompt reduces the cost associated with using an LLM model on Amazon Bedrock, as costs are often based on the

number of tokens processed by the model.
? Token Reduction Strategy:
? Why Option B is Correct:
? Why Other Options are Incorrect:

NEW QUESTION 97
Which option is a benefit of using Amazon SageMaker Model Cards to document AI models?

A. Providing a visually appealing summary of a model's capabilities.
B. Standardizing information about a model's purpose, performance, and limitations.
C. Reducing the overall computational requirements of a model.
D. Physically storing models for archival purposes.

Answer: B

Explanation:
Amazon SageMaker Model Cards provide a standardized way to document important details about an AI model, such as its purpose, performance, intended usage, and known limitations. This enables transparency and compliance while fostering better communication between stakeholders. It does not store models physically or optimize computational requirements. References: AWS SageMaker Model Cards Documentation.

NEW QUESTION 98
A company is training a foundation model (FM). The company wants to increase the accuracy of the model up to a specific acceptance level.
Which solution will meet these requirements?

A. Decrease the batch size.
B. Increase the epochs.
C. Decrease the epochs.
D. Increase the temperature parameter.

Answer: B

Explanation:
Increasing the number of epochs during model training allows the model to learn from the data over more iterations, potentially improving its accuracy up to a certain point. This is a common practice when attempting to reach a specific level of accuracy.
? Option B (Correct): "Increase the epochs": This is the correct answer because
increasing epochs allows the model to learn more from the data, which can lead to higher accuracy.
? Option A: "Decrease the batch size" is incorrect as it mainly affects training speed
and may lead to overfitting but does not directly relate to achieving a specific accuracy level.
? Option C: "Decrease the epochs" is incorrect as it would reduce the training time,
possibly preventing the model from reaching the desired accuracy.
? Option D: "Increase the temperature parameter" is incorrect because temperature affects the randomness of predictions, not model accuracy.
AWS AI Practitioner References:
? Model Training Best Practices on AWS: AWS suggests adjusting training parameters, like the number of epochs, to improve model performance.

NEW QUESTION 101
A company wants to create a chatbot by using a foundation model (FM) on Amazon Bedrock. The FM needs to access encrypted data that is stored in an Amazon S3 bucket.
The data is encrypted with Amazon S3 managed keys (SSE-S3).
The FM encounters a failure when attempting to access the S3 bucket data. Which solution will meet these requirements?

A. Ensure that the role that Amazon Bedrock assumes has permission to decrypt data withthe correct encryption key.
B. Set the access permissions for the S3 buckets to allow public access to enable access over the internet.
C. Use prompt engineering techniques to tell the model to look for information in Amazon S3.
D. Ensure that the S3 data does not contain sensitive information.

Answer: A

Explanation:
Amazon Bedrock needs the appropriate IAM role with permission to access and decrypt data stored in Amazon S3. If the data is encrypted with Amazon S3 managed keys (SSE- S3), the role that Amazon Bedrock assumes must have the required permissions to access and decrypt the encrypted data.
? Option A (Correct): "Ensure that the role that Amazon Bedrock assumes has permission to decrypt data with the correct encryption key": This is the correct solution as it ensures that the AI model can access the encrypted data securely without changing the encryption settings or compromising data security.
? Option B: "Set the access permissions for the S3 buckets to allow public access" is incorrect because it violates security best practices by exposing sensitive data to the public.
? Option C: "Use prompt engineering techniques to tell the model to look for information in Amazon S3" is incorrect as it does not address the encryption and permission issue.
? Option D: "Ensure that the S3 data does not contain sensitive information" is incorrect because it does not solve the access problem related to encryption.
AWS AI Practitioner References:
? Managing Access to Encrypted Data in AWS: AWS recommends using proper IAM roles and policies to control access to encrypted data stored in S3.

NEW QUESTION 103
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

AIF-C01 Practice Exam Features:

* AIF-C01 Questions and Answers Updated Frequently

* AIF-C01 Practice Questions Verified by Expert Senior Certified Staff

* AIF-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AIF-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year