Amazon-Web-Services

Exam Questions AIF-C01

AWS Certified AI Practitioner

NEW QUESTION 1
A company is building a large language model (LLM) question answering chatbot. The company wants to decrease the number of actions call center employees need to take to respond to customer questions.
Which business objective should the company use to evaluate the effect of the LLM chatbot?

A. Website engagement rate
B. Average call duration
C. Corporate social responsibility
D. Regulatory compliance

Answer: B

Explanation:
The business objective to evaluate the effect of an LLM chatbot aimed at reducing the actions required by call center employees should be average call duration.
? Average Call Duration:
? Why Option B is Correct:
? Why Other Options are Incorrect:

NEW QUESTION 2
A law firm wants to build an AI application by using large language models (LLMs). The application will read legal documents and extract key points from the documents. Which solution meets these requirements?

A. Build an automatic named entity recognition system.
B. Create a recommendation engine.
C. Develop a summarization chatbot.
D. Develop a multi-language translation system.

Answer: C

Explanation:
A summarization chatbot is ideal for extracting key points from legal documents. Large language models (LLMs) can be used to summarize complex texts, such as legal documents, making them more accessible and understandable.
? Option C (Correct): "Develop a summarization chatbot": This is the correct answer
because a summarization chatbot uses LLMs to condense and extract key information from text, which is precisely the requirement for reading and summarizing legal documents.
? Option A: "Build an automatic named entity recognition system" is incorrect
because it focuses on identifying specific entities, not summarizing documents.
? Option B: "Create a recommendation engine" is incorrect as it is used to suggest products or content, not summarize text.
? Option D: "Develop a multi-language translation system" is incorrect because translation is unrelated to summarizing text.
AWS AI Practitioner References:
? Using LLMs for Text Summarization on AWS: AWS supports developing summarization tools using its AI services, including Amazon Bedrock.

NEW QUESTION 3
A company wants to make a chatbot to help customers. The chatbot will help solve technical problems without human intervention. The company chose a foundation model (FM) for the chatbot. The chatbot needs to produce responses that adhere to company tone.
Which solution meets these requirements?

A. Set a low limit on the number of tokens the FM can produce.
B. Use batch inferencing to process detailed responses.
C. Experiment and refine the prompt until the FM produces the desired responses.
D. Define a higher number for the temperature parameter.

Answer: C

Explanation:
Experimenting and refining the prompt is the best approach to ensure that the chatbot using a foundation model (FM) produces responses that adhere to the company's tone.
? Prompt Engineering:
? Why Option C is Correct:
? Why Other Options are Incorrect:

NEW QUESTION 4
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 5
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 6
Which term describes the numerical representations of real-world objects and concepts that AI and natural language processing (NLP) models use to improve understanding of textual information?

A. Embeddings
B. Tokens
C. Models
D. Binaries

Answer: A

Explanation:
Embeddings are numerical representations of objects (such as words, sentences, or documents) that capture the objects' semantic meanings in a form that AI and NLP models can easily understand. These representations help models improve their understanding of textual information by representing concepts in a continuous vector space.
? Option A (Correct): "Embeddings": This is the correct term, as embeddings provide
a way for models to learn relationships between different objects in their input space, improving their understanding and processing capabilities.
? Option B: "Tokens" are pieces of text used in processing, but they do not capture
semantic meanings like embeddings do.
? Option C: "Models" are the algorithms that use embeddings and other inputs, not the representations themselves.
? Option D: "Binaries" refer to data represented in binary form, which is unrelated to the concept of embeddings.
AWS AI Practitioner References:
? Understanding Embeddings in AI and NLP: AWS provides resources and tools, like Amazon SageMaker, that utilize embeddings to represent data in formats suitable for machine learning models.

NEW QUESTION 7
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 8
Which AWS feature records details about ML instance data for governance and reporting?

A. Amazon SageMaker Model Cards
B. Amazon SageMaker Debugger
C. Amazon SageMaker Model Monitor
D. Amazon SageMaker JumpStart

Answer: A

Explanation:

Amazon SageMaker Model Cards provide a centralized and standardized repository for documenting machine learning models. They capture key details such as the model's intended use, training and evaluation datasets, performance metrics, ethical considerations, and other relevant information. This documentation facilitates governance and reporting by ensuring that all stakeholders have access to consistent and comprehensive information about each model. While Amazon SageMaker Debugger is used for real-time debugging and monitoring during training, and Amazon SageMaker Model Monitor tracks deployed models for data and prediction quality, neither offers the comprehensive documentation capabilities of Model Cards. Amazon SageMaker JumpStart provides pre-built models and solutions but does not focus on governance documentation.
Reference: Amazon SageMaker Model Cards

NEW QUESTION 9

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 10

A company wants to develop a large language model (LLM) application by using Amazon Bedrock and customer data that is uploaded to Amazon S3. The company's security policy states that each team can access data for only the team's own customers.
Which solution will meet these requirements?

A. Create an Amazon Bedrock custom service role for each team that has access to only the team's customer data.
B. Create a custom service role that has Amazon S3 acces
C. Ask teams to specify the customer name on each Amazon Bedrock request.
D. Redact personal data in Amazon S3. Update the S3 bucket policy to allow team access to customer data.
E. Create one Amazon Bedrock role that has full Amazon S3 acces
F. Create IAM roles for each team that have access to only each team's customer folders.

Answer: A

Explanation:
To comply with the company's security policy, which restricts each team to access data for only their own customers, creating an Amazon Bedrock custom service role for each team is the correct solution.
? Custom Service Role Per Team:
? Why Option A is Correct:
? Why Other Options are Incorrect:
Thus, A is the correct answer to meet the company's security requirements.

NEW QUESTION 10

A company deployed an AI/ML solution to help customer service agents respond to frequently asked questions. The questions can change over time. The company wants to give customer service agents the ability to ask questions and receive automatically generated answers to common customer questions. Which strategy will meet these requirements MOST cost-effectively?

A. Fine-tune the model regularly.
B. Train the model by using context data.
C. Pre-train and benchmark the model by using context data.
D. Use Retrieval Augmented Generation (RAG) with prompt engineering techniques.

Answer: D

Explanation:
 RAG combines large pre-trained models with retrieval mechanisms to fetch relevant context from a knowledge base. This approach is cost-effective as it eliminates the need for frequent model retraining while ensuring responses are contextually accurate and up to date. References: AWS RAG Techniques.

NEW QUESTION 14

A company has thousands of customer support interactions per day and wants to analyze these interactions to identify frequently asked questions and develop insights.
Which AWS service can the company use to meet this requirement?

A. Amazon Lex
B. Amazon Comprehend
C. Amazon Transcribe
D. Amazon Translate

Answer: B

Explanation:
Amazon Comprehend is the correct service to analyze customer support interactions and identify frequently asked questions and insights.
? Amazon Comprehend:

? Why Option B is Correct:
? Why Other Options are Incorrect:

NEW QUESTION 19
A company has petabytes of unlabeled customer data to use for an advertisement campaign. The company wants to classify its customers into tiers to advertise and promote the company's products.
Which methodology should the company use to meet these requirements?

A. Supervised learning
B. Unsupervised learning
C. Reinforcement learning
D. Reinforcement learning from human feedback (RLHF)

Answer: B

Explanation:
Unsupervised learning is the correct methodology for classifying customers into tiers when the data is unlabeled, as it does not require predefined labels or outputs.
? Unsupervised Learning:
? Why Option B is Correct:
? Why Other Options are Incorrect:

NEW QUESTION 20
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

Answer: A

Explanation:
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.

NEW QUESTION 21
A company has terabytes of data in a database that the company can use for business analysis. The company wants to build an AI-based application that can build a SQL query from input text that employees provide. The employees have minimal experience with technology.
Which solution meets these requirements?

A. Generative pre-trained transformers (GPT)
B. Residual neural network
C. Support vector machine
D. WaveNet

Answer: A

Explanation:
Generative Pre-trained Transformers (GPT) are suitable for building an AI-based application that can generate SQL queries from natural language input provided by employees.
? GPT for Natural Language Processing:
? Why Option A is Correct:
? Why Other Options are Incorrect:

NEW QUESTION 25
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 30
A company wants to develop an educational game where users answer questions such as the following: "A jar contains six red, four green, and three yellow marbles. What is the probability of choosing a green marble from the jar?"
Which solution meets these requirements with the LEAST operational overhead?

A. Use supervised learning to create a regression model that will predict probability.
B. Use reinforcement learning to train a model to return the probability.
C. Use code that will calculate probability by using simple rules and computations.
D. Use unsupervised learning to create a model that will estimate probability density.

Answer: C

Explanation:
The problem involves a simple probability calculation that can be handled efficiently by straightforward mathematical rules and computations. Using machine learning techniques would introduce unnecessary complexity and operational overhead.
? Option C (Correct): "Use code that will calculate probability by using simple rules and computations": This is the correct answer because it directly solves the problem with minimal overhead, using basic probability rules.
? Option A: "Use supervised learning to create a regression model" is incorrect as it overcomplicates the solution for a simple probability problem.
? Option B: "Use reinforcement learning to train a model" is incorrect because reinforcement learning is not needed for a simple probability calculation.
? Option D: "Use unsupervised learning to create a model" is incorrect as unsupervised learning is not applicable to this task.
AWS AI Practitioner References:
? Choosing the Right Solution for AI Tasks: AWS recommends using the simplest and most efficient approach to solve a given problem, avoiding unnecessary machine learning techniques for straightforward tasks.

NEW QUESTION 31
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 33
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 an AI model responsibly and minimize bias, it is essential to ensure fairness and transparency throughout the model development and deployment process. This involves detecting and mitigating data imbalances and thoroughly evaluating the model's behavior to understand its impact on different groups.
? Option A (Correct): "Detect imbalances or disparities in the data": This is correct because identifying and addressing data imbalances or disparities is a critical step in reducing bias. AWS provides tools like Amazon SageMaker Clarify to detect bias during data preprocessing and model training.
? Option C (Correct): "Evaluate the model's behavior so that the company can provide transparency to stakeholders": This is correct because evaluating the model's behavior for fairness and accuracy is key to ensuring that stakeholders understand how the model makes decisions. Transparency is a crucial aspect of responsible AI.
? Option B: "Ensure that the model runs frequently" is incorrect because the frequency of model runs does not address bias.
? Option D: "Use the Recall-Oriented Understudy for Gisting Evaluation (ROUGE) technique to ensure that the model is 100% accurate" is incorrect because ROUGE is a metric for evaluating the quality of text summarization models, not for minimizing bias.
? Option E: "Ensure that the model's inference time is within the accepted limits" is incorrect as it relates to performance, not bias reduction.
AWS AI Practitioner References:
? Amazon SageMaker Clarify: AWS offers tools such as SageMaker Clarify for detecting bias in datasets and models, and for understanding model behavior to ensure fairness and transparency.
? Responsible AI Practices: AWS promotes responsible AI by advocating for fairness, transparency, and inclusivity in model development and deployment.

NEW QUESTION 37
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 41
A company makes forecasts each quarter to decide how to optimize operations to meet expected demand. The company uses ML models to make these forecasts.
An AI practitioner is writing a report about the trained ML models to provide transparency and explainability to company stakeholders.
What should the AI practitioner include in the report to meet the transparency and explainability requirements?

A. Code for model training
B. Partial dependence plots (PDPs)
C. Sample data for training
D. Model convergence tables

Answer: B

Explanation:
Partial dependence plots (PDPs) are visual tools used to show the relationship between a feature (or a set of features) in the data and the predicted outcome of a machine learning model. They are highly effective for providing transparency and explainability of the model's behavior to stakeholders by illustrating how different input variables impact the model's predictions.
? Option B (Correct): "Partial dependence plots (PDPs)": This is the correct answer because PDPs help to interpret how the model's predictions change with varying values of input features, providing stakeholders with a clearer understanding of the model's decision-making process.
? Option A: "Code for model training" is incorrect because providing the raw code for model training may not offer transparency or explainability to non-technical stakeholders.
? Option C: "Sample data for training" is incorrect as sample data alone does not explain how the model works or its decision-making process.
? Option D: "Model convergence tables" is incorrect. While convergence tables can show the training process, they do not provide insights into how input features affect the model's predictions.
AWS AI Practitioner References:
? Explainability in AWS Machine Learning: AWS provides various tools for model explainability, such as Amazon SageMaker Clarify, which includes PDPs to help explain the impact of different features on the model??s predictions.

NEW QUESTION 44
......

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

100% Actual & Verified — Instant Download, Please Click
Order The AIF-C01 Practice Test Here