Pass4Leader

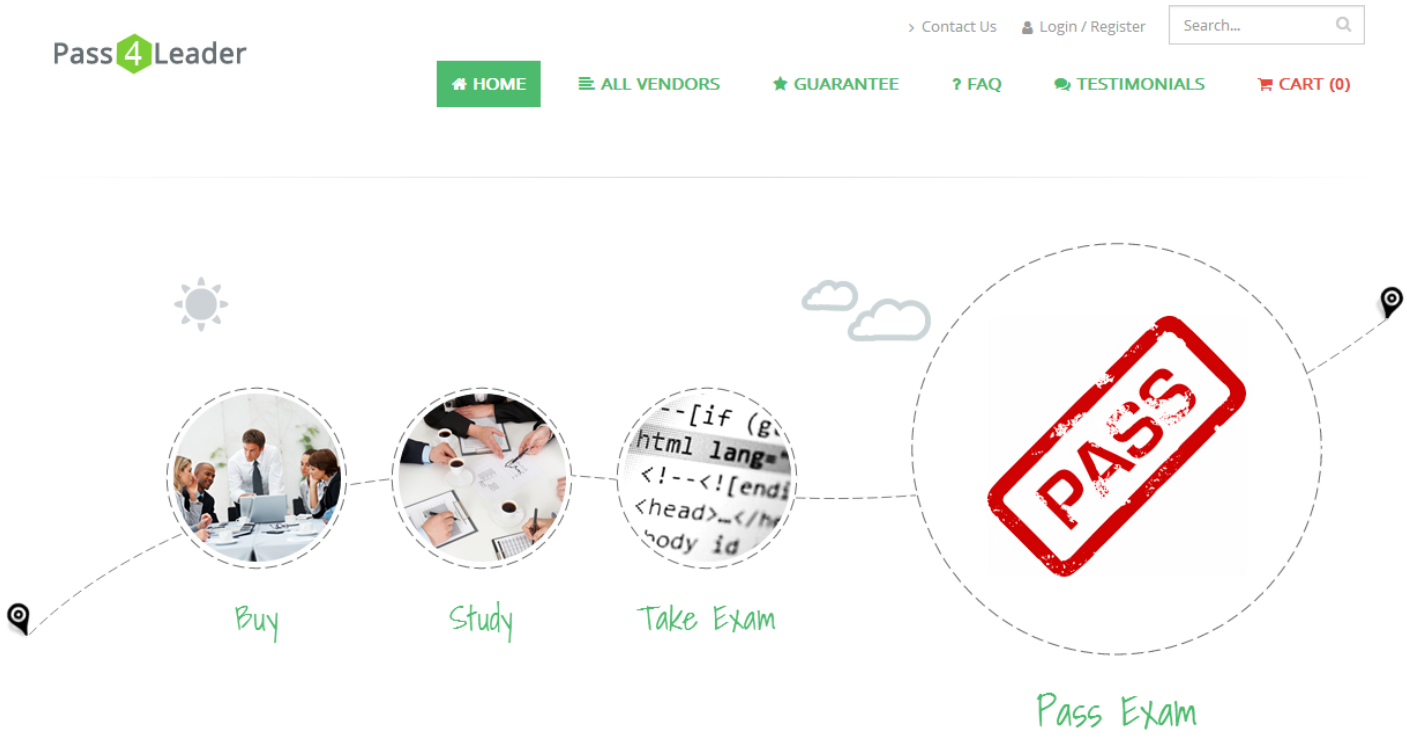

Try **Online Engine** before you buy

We're not the only ones **excited** about Pass4Leader Practice Material ...

66966+ customers in 100+ countries use Pass4Leader Test Engine. Meet our customers.

JET ORANGE iCompany Paradoxx iMessenger VOREED GetCustom

http://www.pass4leader.com/

Latest Exam Guide & Learning Materials

Exam : **AIF-C01-KR**

Title : AWS Certified AI Practitioner (AIF-C01 Korean Version)

Vendor : Amazon

Version : DEMO

QUESTION NO: 1
한 소프트웨어 회사가 워크플로 자동화에 대규모 언어 모델(LLM)을 사용하려고 합니다. 이 애플리케이션은 사용자 메시지를 JSON 파일로 변환합니다. 회사는 이 JSON 파일을 데이터 파이프라인의 입력으로 사용합니다.
해당 회사에는 사용자 메시지와 출력 JSON 파일이 포함된 레이블이 지정된 데이터 세트가 있습니다.
어떤 솔루션이 LLM의 워크플로 자동화를 훈련시킬 수 있을까요?
A. 비지도 학습
B. 사전 훈련 계속
C. 미세 조정
D. 인간 피드백을 통한 강화 학습(RLHF)
Answer: C
Explanation:
Fine-tuning is the process of training a pre-trained LLM with a labeled dataset specific to a desired task-in this case, mapping user messages to JSON outputs. Fine-tuning leverages supervised learning to specialize the model's outputs.
C is correct:
"Fine-tuning is a supervised learning approach in which a model is further trained on a custom, labeled dataset to adapt to a specific use case." (Reference: Amazon Bedrock Fine-Tuning, AWS Certified AI Practitioner Study Guide) A is incorrect-unsupervised learning does not use labeled data.
B (continued pre-training) uses unlabeled data.
D (RLHF) uses reward signals and human feedback, not direct labeled input/output pairs.

QUESTION NO: 2
AI와 자연어 처리(NLP) 모델이 텍스트 정보에 대한 이해를 높이기 위해 사용하는 실제 세계 객체와 개념의 수치적 표현을 설명하는 용어는 무엇입니까?
A. 임베딩
B. 토큰
C. 모델
D. 바이너리
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.

QUESTION NO: 3
한 회사에서는 Amazon Bedrock Agents를 사용하여 비즈니스 워크플로를 자동화하는 애플리케이션을 구축하고 있습니다.
A. 시각, 오디오 및 텍스트 입력을 처리하기 위해 기초 모델(FM)을 호출합니다.
B. 촉진 전략을 통해 기초 모델(FM)을 강화합니다.
C. 사용자에게 외부 데이터 소스 및 API 쿼리에 대한 전체 제어권을 제공합니다.
D. 사용자 입력을 평가하고 여러 작업에 대한 작업을 조율합니다.
Answer: D
Explanation:
The correct answer is D. Amazon Bedrock Agents are used to orchestrate and execute complex workflows by connecting foundation models with APIs, databases, and tools. According to AWS documentation, agents interpret user inputs, plan the necessary steps, call external APIs or systems, and return structured results. This allows the model to go beyond text generation into full automation workflows-such as booking tasks, querying internal systems, or summarizing reports. Option A describes multi-modal models, B refers to prompt tuning, and C misstates control delegation; agents act autonomously based on model reasoning. Thus, Bedrock Agents function as intelligent orchestrators, handling multi-step task execution through integrated tool use.
Referenced AWS AI/ML Documents and Study Guides:
Amazon Bedrock Developer Guide - Agents Overview
AWS Generative AI Best Practices - Workflow Orchestration

QUESTION NO: 4
한 회사가 Amazon Bedrock에서 제3자 모델을 사용하여 기밀 문서를 분석합니다. 이 회사는 데이터 개인정보 보호에 대해 우려하고 있습니다. Amazon Bedrock이 데이터 개인정보를 보호하는 방식을 설명하는 다음 문장은 무엇입니까?
A. 사용자 입력과 모델 출력은 익명화되어 타사 모델 제공자와 공유됩니다.
B. 사용자 입력 및 모델 출력은 제3자 모델 공급자와 공유되지 않습니다.
C. 사용자 입력은 비밀로 유지되지만 모델 출력은 제3자 모델 공급자와 공유됩니다.
D. 사용자 입력과 모델 출력은 입력과 출력이 제3자 모델 공급자와 공유되기 전에 삭제됩니다.
Answer: B
Explanation:
Comprehensive and Detailed Explanation from AWS AI Documents:
Amazon Bedrock ensures data privacy and security by not sharing customer inputs or outputs with third-party model providers.
The models are accessed via Bedrock's API isolation layer, meaning that model providers do not see your data.
Customer data is not used for training or improving foundation models unless customers

explicitly opt in.
From AWS Docs:
"Amazon Bedrock does not share your inputs and outputs with third-party model providers. Your data remains private, and is not used to improve the foundation models." This ensures full data privacy, especially for sensitive use cases like confidential documents.
Reference:
AWS Documentation - Data privacy in Amazon Bedrock

QUESTION NO: 5
한 회사가 Amazon Bedrock에서 감정 분석을 위해 대규모 언어 모델(LLM)을 사용하고자 합니다. 이 회사는 하나의 프롬프트에 얼마나 많은 정보가 들어갈 수 있는지 알고 싶어합니다. 어떤 고려사항이 회사의 결정에 영향을 미칠까요?
A. 온도
B. 컨텍스트 창
C. 배치 크기
D. 모델 크기
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:
The context window is the maximum amount of text (measured in tokens) that a language model can process in a single pass.
For LLM applications, the size of the context window limits how much input data, such as text for sentiment analysis, can be fed into the model at once.
Why Option B is Correct:
Determines Prompt Size: The context window size directly informs how much information (e.g., words or sentences) can fit in one prompt.
Model Capacity: The larger the context window, the more information the model can consider for generating outputs.
Why Other Options are Incorrect:
A). Temperature: Controls randomness in model outputs but does not affect the prompt size.
C). Batch size: Refers to the number of training samples processed in one iteration, not the amount of information in a prompt.
D). Model size: Refers to the number of parameters in the model, not the input size for a single prompt.

QUESTION NO: 6
핫스팟
다음 목록에서 각 문장에 맞는 AI 용어를 선택하세요. 각 AI 용어는 한 번씩만 선택해야 합니다. (세 개를 선택하세요.)
* 인공지능
* 딥러닝
* 밀

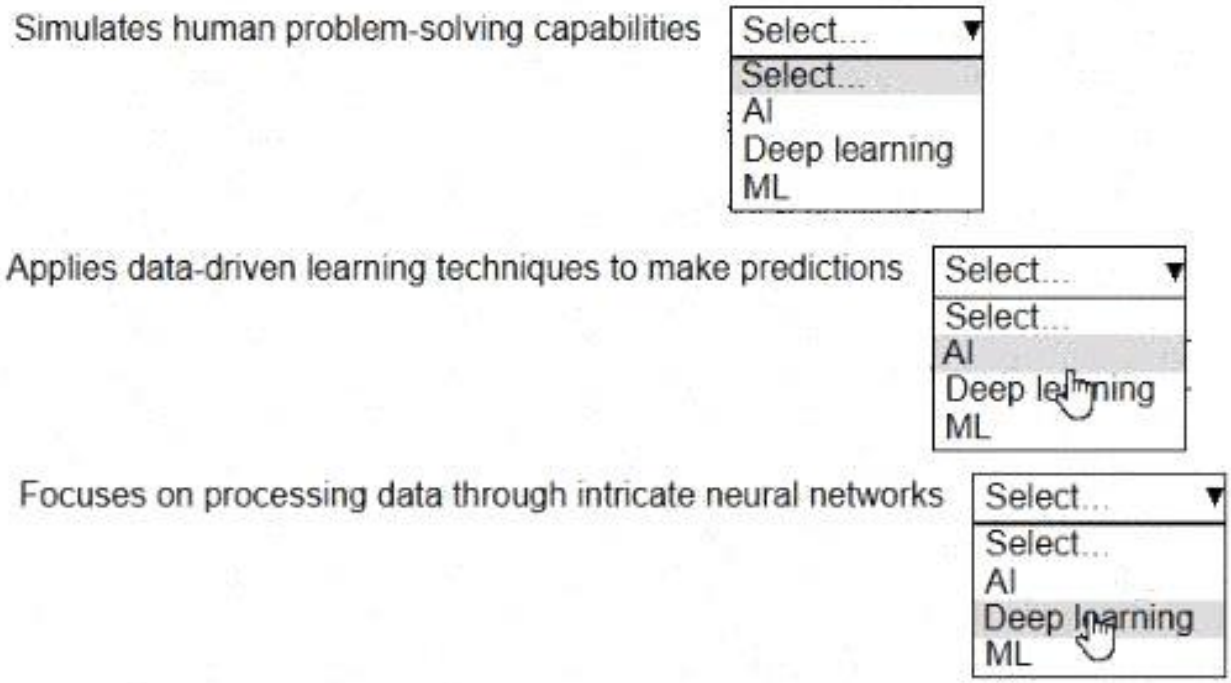

Answer:

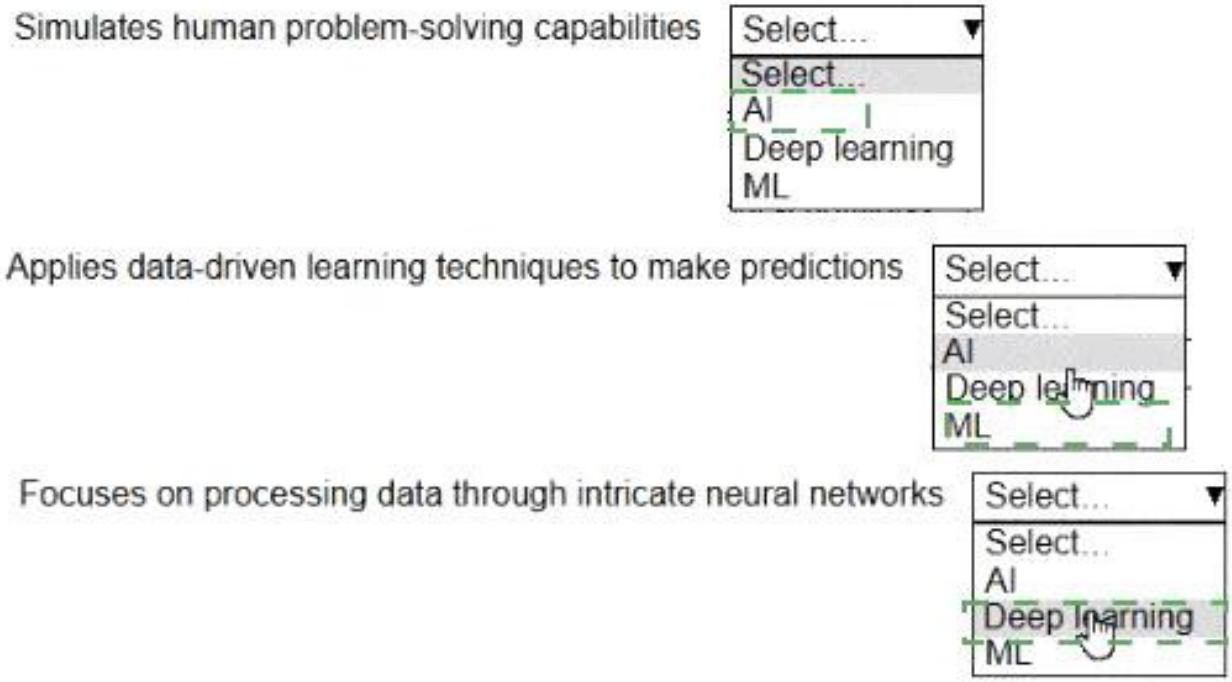

Explanation:
Artificial Intelligence (AI) is the broad field focused on simulating human problem-solving and cognitive abilities, including reasoning, perception, and decision-making.
(Reference: AWS Certified AI Practitioner Official Study Guide)
Machine Learning (ML) is a subset of AI that uses data-driven algorithms to identify patterns and make predictions without explicit programming for each specific task.
(Reference: AWS Machine Learning Overview)
Deep learning is a subset of ML that uses neural networks with many layers (deep neural networks) to process complex data and extract high-level features.

(Reference: AWS Deep Learning on AWS)

QUESTION NO: 7
한 회사가 휴대폰 카메라를 사용하여 곤충 물림을 진단하고 치료하는 모바일 머신러닝 앱을 개발하고 있습니다. 이 회사는 전 세계 다양한 성별, 민족, 지리적 위치에서 수집된 곤충 물림 사진 데이터 세트를 사용하여 이미지 분류 모델을 학습시키고자 합니다.
이 시나리오에서 회사는 어떤 책임 있는 AI 원칙을 보여주고 있나요?
A. 공정성
B. 설명 가능성
C. 거버넌스
D. 투명성
Answer: A
Explanation:
The company is training an image classification model for diagnosing insect bites using a diverse dataset that includes photos from different genders, ethnicities, and geographic locations. This approach demonstrates the principle of fairness in responsible AI, as it aims to reduce bias and ensure the model performs equitably across diverse populations.
Exact Extract from AWS AI Documents:
From the AWS AI Practitioner Learning Path:
" Fairness in AI involves ensuring that models do not exhibit bias against certain groups and perform equitably across diverse populations. This can be achieved by training models on diverse datasets that represent various demographics, such as gender, ethnicity, and geographic location. " (Source: AWS AI Practitioner Learning Path, Module on Responsible AI) Detailed Explanation:
Option A: FairnessThis is the correct answer. By using a diverse dataset, the company ensures the model is less likely to be biased against specific groups, promoting fairness in its predictions and treatments for insect bites.
Option B: ExplainabilityExplainability refers to making the model's decisions understandable to users, such as byproviding insights into how predictions are made. The scenario focuses on dataset diversity, not explainability.
Option C: GovernanceGovernance involves establishing policies and processes to manage AI systems, such as compliance and oversight. The scenario does not describe governance mechanisms.
Option D: TransparencyTransparency involves disclosing how a model works, its limitations, and its data sources. While transparency is important, the scenario specifically highlights the diversity of the dataset, which aligns more directly with fairness.
References:
AWS AI Practitioner Learning Path: Module on Responsible AI
AWS Documentation: Responsible AI Principles (https://aws.amazon.com/machine-learning/responsible-ai/) Amazon SageMaker Developer Guide: Bias and Fairness in ML (https://docs.aws.amazon.com/sagemaker
/latest/dg/clarify-bias.html)

QUESTION NO: 8
한 회사가 공개적으로 이용 가능한 파운데이션 모델(FM)에 대한 직원의 접근을 제어하려고

합니다. 이러한 요구 사항을 충족하는 솔루션은 무엇입니까?
A. AWS Cost Explorer에서 비용 및 사용량 보고서를 분석합니다.
B. AWS Artifact에서 AWS 보안 및 규정 준수 문서를 다운로드합니다.
C. 검색 가능한 FM을 제한하도록 Amazon SageMaker JumpStart를 구성합니다.
D. Amazon OpenSearch Service를 사용하여 하이브리드 검색 솔루션을 구축합니다.
Answer: C
Explanation:
The correct answer is C because Amazon SageMaker JumpStart provides administrative controls that allow organizations to manage and restrict access to foundation models within the AWS environment.
According to the official AWS documentation:
" Amazon SageMaker JumpStart provides model access management capabilities that enable administrators to control which foundation models are visible and usable by end users. Using AWS Identity and Access Management (IAM) policies, you can restrict access to specific models or completely disable model discovery in JumpStart. " This allows companies to enforce governance over which FMs their users can see and interact with, satisfying the requirement to control employee access to publicly available foundation models.
Explanation of other options:
A). AWS Cost Explorer is used to analyze billing and usage data but does not control access to services or models. It is helpful for budgeting and visibility, not access control.
B). AWS Artifact provides access to compliance reports and certifications, not tools for controlling user access to ML models.
D). Amazon OpenSearch Service is used for search and analytics on structured and unstructured data. It does not provide access control mechanisms for foundation models.
Referenced AWS AI/ML Documents and Study Guides:
Amazon SageMaker JumpStart Documentation - Model Access Management
AWS IAM Documentation - Restricting Access to SageMaker Resources
AWS Machine Learning Specialty Certification Guide - Security and Governance Section

QUESTION NO: 9
한 회사가 특정 품목의 가격을 예측하는 새로운 모델을 개발하고 있습니다. 이 모델은 훈련 데이터 세트에서 좋은 성과를 보였습니다. 회사가 모델을 프로덕션에 배포했을 때 모델의 성과가 상당히 떨어졌습니다.
이 문제를 완화하기 위해 회사는 무엇을 해야 할까?
A. 학습에 사용되는 데이터 양을 줄입니다.
B. 모델에 하이퍼파라미터를 추가합니다.
C. 훈련에 사용되는 데이터의 양을 늘립니다.
D. 모델 학습 시간을 늘립니다.
Answer: C
Explanation:
When a model performs well on the training data but poorly in production, it is often due to overfitting.
Overfitting occurs when a model learns patterns and noise specific to the training data, which

does not generalize well to new, unseen data in production. Increasing the volume of data used in training can help mitigate this problem by providing a more diverse and representative dataset, which helps the model generalize better.
Option C (Correct): " Increase the volume of data that is used in training " : Increasing the data volume can help the model learn more generalized patterns rather than specific features of the training dataset, reducing overfitting and improving performance in production.
Option A: " Reduce the volume of data that is used in training " is incorrect, as reducing data volume would likely worsen the overfitting problem.
Option B: " Add hyperparameters to the model " is incorrect because adding hyperparameters alone does not address the issue of data diversity or model generalization.
Option D: " Increase the model training time " is incorrect because simply increasing training time does not prevent overfitting; the model needs more diverse data.
AWS AI Practitioner References:
Best Practices for Model Training on AWS: AWS recommends using a larger and more diverse training dataset to improve a model ' s generalization capability and reduce the risk of overfitting.

QUESTION NO: 10
어떤 AWS 서비스나 기능이 AI 개발 팀이 팀의 VPC 내에서 기반 모델(FM)을 신속하게 배포하고 사용하는 데 도움이 될 수 있습니까?
A. Amazon 개인화
B. Amazon SageMaker JumpStart
C. Amazon Bedrock Playground인 PartyRock
D. Amazon SageMaker 엔드포인트
Answer: B
Explanation:
Amazon SageMaker JumpStart is the correct service for quickly deploying and consuming a foundation model (FM) within a team ' s VPC.
Amazon SageMaker JumpStart:
Provides access to a wide range of pre-trained models and solutions that can be easily deployed and consumed within a VPC.
Designed to simplify and accelerate the deployment of machine learning models, including foundation models.
Why Option B is Correct:
Rapid Deployment: JumpStart allows for quick deployment of models with minimal configuration, directly within a secure VPC environment.
Ease of Use: Provides a user-friendly interface to select and deploy models, reducing the time to value.
Why Other Options are Incorrect:
A). Amazon Personalize: Focuses on creating personalized recommendations, not deploying foundation models.
C). PartyRock: Not a recognized AWS service.
D). Amazon SageMaker endpoints: Endpoints are for deploying specific models, not a feature for quickly starting with pre-trained foundation models.

QUESTION NO: 11
한 회사가 Amazon Nova Canvas 모델을 사용하여 이미지를 생성하고 있습니다. 이 모델은 이미지를 성공적으로 생성합니다.
회사에서는 모델이 생성된 이미지에 특정 항목을 포함하는 것을 방지해야 합니다.
어떤 솔루션이 이 요구 사항을 충족할 수 있을까요?
A. 더 높은 온도 값을 사용합니다.
B. 더 자세한 프롬프트를 사용하세요.
C. 부정적인 프롬프트를 사용합니다.
D. 다른 기초 모델(FM)을 사용합니다.
Answer: C
Explanation:
The correct answer is C - Use a negative prompt. Negative prompts instruct a generative image model to avoid certain features, objects, or styles in the output. This technique is fully supported by models like Amazon Nova Canvas on Bedrock, which are based on diffusion or image generation architectures.
According to AWS documentation, negative prompts refine output control by telling the model what not to include, thereby improving brand alignment, compliance, or creative direction. A higher temperature increases randomness, not control. A detailed prompt helps, but without exclusion instructions, the model may still include unwanted elements. Changing the model may yield better output but doesn't directly solve this control requirement. Negative prompts are purpose-built for this scenario.
Referenced AWS AI/ML Documents and Study Guides:
Amazon Bedrock Documentation - Prompt Engineering for Image Models
AWS Generative AI Guide - Controlled Generation with Negative Prompts

QUESTION NO: 12
Amazon SageMaker Clarify는 어떤 기능을 제공합니까?
A. RAG(Retrieval Augmented Generation) 워크플로를 통합합니다.
B. 프로덕션에서 ML 모델의 품질을 모니터링합니다.
C. ML 모델에 대한 중요한 세부 정보를 문서화합니다.
D. 데이터 준비 중 잠재적인 편향을 식별합니다.
Answer: D
Explanation:
Exploratory data analysis (EDA) involves understanding the data by visualizing it, calculating statistics, and creating correlation matrices. This stage helps identify patterns, relationships, and anomalies in the data, which can guide further steps in the ML pipeline.
Option C (Correct): " Exploratory data analysis " : This is the correct answer as the tasks described (correlation matrix, calculating statistics, visualizing data) are all part of the EDA process.
Option A: " Data pre-processing " is incorrect because it involves cleaning and transforming data, not initial analysis.
Option B: " Feature engineering " is incorrect because it involves creating new features from raw data, not analyzing the data ' s existing structure.
Option D: " Hyperparameter tuning " is incorrect because it refers to optimizing model

parameters, not analyzing the data.
AWS AI Practitioner References:
Stages of the Machine Learning Pipeline: AWS outlines EDA as the initial phase of understanding and exploring data before moving to more specific preprocessing, feature engineering, and model training stages.

QUESTION NO: 13
한 회사는 Generative AI Security Scoping Matrix를 사용하여 솔루션에 대한 보안 책임을 평가하고 있습니다. 이 회사는 매트릭스를 기반으로 4가지 다른 솔루션 범위를 식별했습니다. 어떤 솔루션 범위가 회사에 보안 책임에 대한 가장 큰 소유권을 제공합니까?
A. 생성 AI 기능이 내장된 타사 엔터프라이즈 애플리케이션을 사용합니다.
B. 기존 타사 생성 AI 기반 모델(FM)을 사용하여 애플리케이션을 빌드합니다.
C. 비즈니스에 특화된 데이터를 사용하여 모델을 미세 조정하여 기존 타사 생성 AI 기반 모델(FM)을 개선합니다.
D. 고객이 소유한 특정 데이터를 사용하여 처음부터 생성 AI 모델을 구축하고 훈련합니다.
Answer: D
Explanation:
Building and training a generative AI model from scratch provides the company with the most ownership and control over security responsibilities. In this scenario, the company is responsible for all aspects of the security of the data, the model, and the infrastructure.
Option D (Correct): " Building and training a generative AI model from scratch by using specific data that a customer owns " : This is the correct answer because it involves complete ownership of the model, data, and infrastructure, giving the company the highest level of responsibility for security.
Option A: " Using a third-party enterprise application that has embedded generative AI features " is incorrect as the company has minimal control over the security of the AI features embedded within a third-party application.
Option B: " Building an application using an existing third-party generative AI foundation model (FM) " is incorrect because security responsibilities are shared with the third-party model provider.
Option C: " Refining an existing third-party generative AI FM by fine-tuning the model with business-specific data " is incorrect as the foundation model and part of the security responsibilities are still managed by the third party.
AWS AI Practitioner References:
Generative AI Security Scoping Matrix on AWS: AWS provides a security responsibility matrix that outlines varying levels of control and responsibility depending on the approach to developing and using AI models.

QUESTION NO: 14
한 회사가 직원들이 내부 데이터를 조회할 수 있도록 인공지능 비서를 개발하려고 합니다. 어떤 AWS 서비스가 이 요구 사항을 충족할까요?
A. 아마존 레코그니션
B. 아마존 텍스트랙트
C. 아마존 렉스
D. 아마존 Q 비즈니스

Answer: C

QUESTION NO: 15
한 병원에서 환자에게 맞춤형 치료 추천을 제공하는 AI 시스템을 개발했습니다. AI 시스템은 추천의 근거를 제시하고, 의사와 환자가 그 통찰력을 활용할 수 있도록 해야 합니다.
A. 설명 가능성
B. 개인정보 보호 및 보안
C. 공정성
D. 데이터 거버넌스
Answer: A
Explanation:
The correct answer is A - Explainability. According to AWS Responsible AI documentation, explainability refers to an AI system's ability to clearly communicate why it produced a given result. In healthcare, clinical decision support systems must provide traceable, understandable reasoning, especially when generating treatment recommendations. AWS highlights explainability as critical for high-impact domains such as medicine because doctors and patients must trust and understand the basis for AI-driven decisions. Tools like Amazon SageMaker Clarify support feature attribution, helping clinicians understand which patient factors (e.
g., age, symptoms, lab values) influenced a recommendation. Privacy and security (B) protect data but do not provide rationale. Fairness (C) ensures equitable treatment across demographics but does not explain decisions. Data governance (D) focuses on handling and controlling data, not model decision transparency.
Explainability is the AWS principle that ensures clinical users can interpret, validate, and rely on AI- generated recommendations for patient care.
Referenced AWS Documentation:
AWS Responsible AI Whitepaper - Explainability
Amazon SageMaker Clarify - Feature Attribution and Model Insights

QUESTION NO: 16
어떤 AWS 서비스가 사용자가 생성적 AI 애플리케이션을 구축하고 확장하는 데 도움이 되는 기초 모델(FM)을 제공합니까?
A. Amazon Q 개발자
B. 아마존 베드락
C. 아마존 켄드라
D. 아마존 컴프리헨드
Answer: B
Explanation:
The correct answer is Amazon Bedrock, AWS's fully managed service for building and scaling generative AI applications using foundation models (FMs). Bedrock gives developers access to models from leading providers such as Anthropic (Claude), Meta (Llama), Mistral, Cohere, and Amazon Titan. Users can invoke these models via API without managing infrastructure or model training. According to AWS documentation, Bedrock supports tasks such as text generation, summarization, question answering, image generation, and RAG

workflows with minimal setup. It supports both on-demand and provisioned throughput modes and integrates with features like Guardrails, Knowledge Bases, and Agents for secure, enterprise-grade applications. Amazon Q Developer is a generative AI tool for developers, but it doesn't host or scale models.
Amazon Kendra is an intelligent search engine, and Amazon Comprehend is used for NLP tasks like entity extraction-not foundation model access.
Referenced AWS AI/ML Documents and Study Guides:
Amazon Bedrock Developer Guide - Foundation Models and Use Cases
AWS Certified ML Specialty Guide - Generative AI on AWS

QUESTION NO: 17
한 회사가 3개월 안에 대규모 언어 모델(LLM)을 개선하여 콘텐츠 검열 기능을 강화하고자 합니다. 이 모델은 회사의 가치관과 윤리에 따라 콘텐츠를 검열해야 하며, 새로운 트렌드와 문제성 콘텐츠 유형에도 대응할 수 있어야 합니다.
어떤 솔루션이 이러한 요구 사항을 충족할까요?
A. 대량의 텍스트 기반 인터넷 콘텐츠에 대한 지속적인 사전 학습을 수행합니다.
B. 과거 검토 결정에 대한 고품질 데이터 세트를 생성합니다.
C. 다양한 출처의 일반적인 윤리 지침 세트를 기반으로 LLM을 세밀하게 조정합니다.
D. 숙련된 진행자의 실시간 입력을 사용하여 인간 피드백 기반 강화 학습(RLHF)을 수행합니다.
Answer: D
Explanation:
Reinforcement learning from human feedback (RLHF) is the most effective approach to align large language models with company-specific values, ethics, and evolving moderation requirements . AWS documentation explains that RLHF uses direct human input to guide model behavior , enabling models to learn preferences that cannot be fully captured through static datasets or generic fine-tuning.
In this scenario, the company requires improvement within a short time frame of three months , alignment with organizational ethics , and adaptability to emerging trends and new forms of harmful content .
RLHF meets these needs by incorporating real-time feedback from skilled human moderators , allowing the model to rapidly adjust its responses based on expert judgment.
AWS highlights that RLHF is particularly valuable for content moderation, safety alignment, and policy enforcement , where nuanced decisions and evolving standards are common. By rewarding desirable behaviors and penalizing undesirable outputs, the model continuously improves in a controlled and targeted manner.
The other options are less suitable. Continuous pre-training on large internet datasets is time-consuming, resource-intensive, and may introduce content misaligned with company values. Historical moderation datasets may not reflect new or emerging content patterns. Fine-tuning on general ethical guidelines lacks the specificity required for company-defined moderation policies and does not adapt quickly to new risks.
AWS positions RLHF as a key technique in responsible generative AI development , enabling organizations to maintain human oversight while improving model safety and alignment.
Therefore, using RLHF with real- time input from skilled moderators is the most effective and compliant solution for this use case.

QUESTION NO: 18
한 회사에서는 Amazon SageMaker를 사용하여 AI 모델을 개발하고 있습니다.
AI 모델 수명 주기 워크플로의 각 단계에 맞는 SageMaker 기능 또는 리소스를 다음 목록에서 선택하세요. 각 SageMaker 기능 또는 리소스는 한 번만 선택하거나 전혀 선택하지 않아야 합니다. (두 개를 선택하세요.)
* SageMaker Clarify
* SageMaker 모델 레지스트리
* SageMaker 서버리스 추론

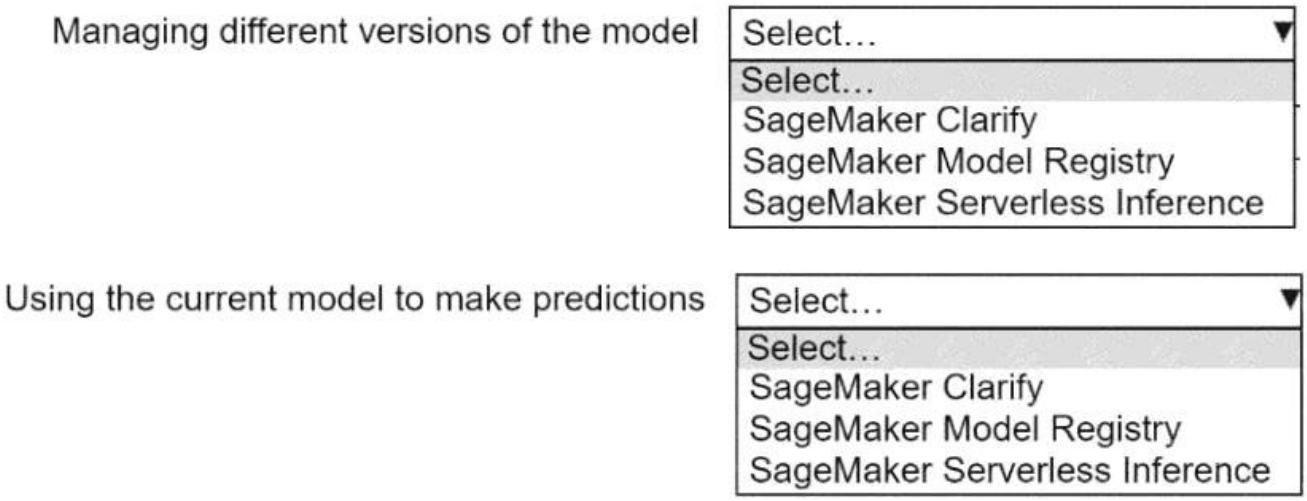

Answer:

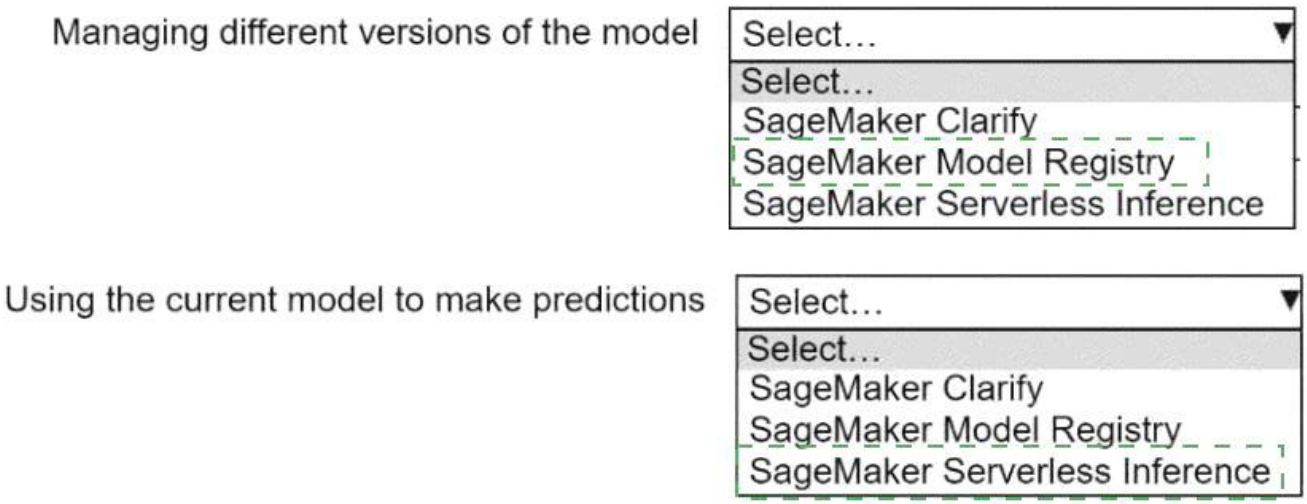

Explanation:
SageMaker Model Registry, SageMaker Serverless interference
This question requires selecting the appropriate Amazon SageMaker feature for two distinct steps in the AI model lifecycle. Let's break down each step and evaluate the options:
Step 1: Managing different versions of the model
The goal here is to identify a SageMaker feature that supports version control and management of machine learning models. Let's analyze the options:
SageMaker Clarify: This feature is used to detect bias in models and explain model predictions, helping with fairness and interpretability. It does not provide functionality for managing model versions.

SageMaker Model Registry: This is a centralized repository in Amazon SageMaker that allows users to catalog, manage, and track different versions of machine learning models. It supports model versioning, approval workflows, and deployment tracking, making it ideal for managing different versions of a model.
SageMaker Serverless Inference: This feature enables users to deploy models for inference without managing servers, automatically scaling based on demand. It is focused on inference (predictions), not on managing model versions.
Conclusion for Step 1: The SageMaker Model Registry is the correct choice for managing different versions of the model.
Exact Extract Reference: According to the AWS SageMaker documentation, "The SageMaker Model Registry allows you to catalog models for production, manage model versions, associate metadata, and manage approval status for deployment." (Source: AWS SageMaker Documentation - Model Registry, https://docs.aws.amazon.com/sagemaker/latest/dg/model-registry.html).
Step 2: Using the current model to make predictions
The goal here is to identify a SageMaker feature that facilitates making predictions (inference) with a deployed model. Let's evaluate the options:
SageMaker Clarify: As mentioned, this feature focuses on bias detection and explainability, not on performing inference or making predictions.
SageMaker Model Registry: While the Model Registry helps manage and catalog models, it is not used directly for making predictions. It can store models, but the actual inference process requires a deployment mechanism.
SageMaker Serverless Inference: This feature allows users to deploy models for inference without managing infrastructure. It automatically scales based on traffic and is specifically designed for making predictions in a cost-efficient, serverless manner.
Conclusion for Step 2: SageMaker Serverless Inference is the correct choice for using the current model to make predictions.
Exact Extract Reference: The AWS documentation states, "SageMaker Serverless Inference is a deployment option that allows you to deploy machine learning models for inference without configuring or managing servers. It automatically scales to handle inference requests, making it ideal for workloads with intermittent or unpredictable traffic." (Source: AWS SageMaker Documentation - Serverless Inference, https://docs.aws.amazon.com/sagemaker/latest/dg/serverless-inference.html).
Why Not Use the Same Feature Twice?
The question specifies that each SageMaker feature or resource should be selected one time or not at all. Since SageMaker Model Registry is used for version management and SageMaker Serverless Inference is used for predictions, each feature is selected exactly once. SageMaker Clarify is not applicable to either step, so it is not selected at all, fulfilling the question's requirements.
References:
AWS SageMaker Documentation: Model Registry (https://docs.aws.amazon.com/sagemaker/latest/dg/model- registry.html) AWS SageMaker Documentation: Serverless Inference (https://docs.aws.amazon.com/sagemaker/latest/dg /serverless-inference.html)
AWS AI Practitioner Study Guide (conceptual alignment with SageMaker features for model

lifecycle management and inference) Let's format this question according to the specified structure and provide a detailed, verified answer based on AWS AI Practitioner knowledge and official AWS documentation. The question focuses on selecting an AWS database service that supports storage and queries of embeddings as vectors, which is relevant to generative AI applications.

QUESTION NO: 19
AI 전문가가 Amazon Bedrock에서 LLM(Learning Language Model)을 평가 도구로 사용하여 실제 운영 환경에서 에이전트 응답 품질을 평가하고 있습니다. 이 전문가는 에이전트 응답이 각 질문이나 프롬프트의 모든 부분을 얼마나 철저하게 다루는지 평가하는 내장 지표를 적용하고자 합니다.
어떤 지표가 이러한 요구 사항을 충족할까요?
A. 요약 평가를 위한 회상 중심 보조 연구(ROUGE)
B. 완전성
C. 지시사항을 따르세요
D. 거부
Answer: B
Explanation:
Comprehensive and Detailed Explanation From Exact AWS AI documents:
In Amazon Bedrock evaluations, Completeness measures how thoroughly a model or agent response addresses all aspects of the user prompt.
AWS evaluation guidance for LLM-as-a-judge explains that:
Completeness focuses on coverage of prompt requirements
It is especially useful for evaluating multi-part questions
It is a built-in qualitative metric in agent evaluation workflows
Why the other options are incorrect:
ROUGE (A) measures text overlap, mainly for summarization.
Following instructions (C) evaluates adherence, not coverage.
Refusal (D) measures appropriate refusal behavior.
AWS AI document references:
Amazon Bedrock Model Evaluation
LLM-as-a-Judge Metrics
Evaluating Agent Responses on AWS

QUESTION NO: 20
한 회사가 센서 데이터에서 이상 패턴을 감지하는 ML 모델을 구축하려고 합니다. 하지만 이 회사에는 학습용 레이블이 지정된 데이터가 없습니다. 이러한 요구 사항을 충족하는 ML 방법은 무엇일까요?
A. 선형 회귀
B. 분류
C. 의사결정 트리
D. 자동 인코더
Answer: D
Explanation:
The correct answer is D because autoencoders are an unsupervised machine learning

method commonly used for anomaly detection when labeled data is not available.
From AWS documentation:
" Autoencoders learn to compress and reconstruct input data. During anomaly detection, they learn normal patterns in data. Data points that the model cannot accurately reconstruct are flagged as anomalies. " This approach is ideal when there is no labeled data and when patterns must be learned based on normal behavior alone - a common situation in IoT sensor data environments.
Explanation of other options:
A). Linear regression requires labeled data and is used for predicting continuous values.
B). Classification requires labeled data to assign inputs into categories.
C). Decision trees are supervised learning models and also require labeled datasets.
Referenced AWS AI/ML Documents and Study Guides:
AWS Machine Learning Specialty Guide - Unsupervised Learning Techniques Amazon SageMaker Examples - Anomaly Detection Using Autoencoders

QUESTION NO: 21
ML 라이프사이클의 어느 단계에서 규정 준수 및 규제 요구 사항이 결정됩니까?
A. 기능 엔지니어링
B. 모델 학습
C. 데이터 수집
D. 비즈니스 목표 식별
Answer: D
Explanation:
The business goal identification phase of the ML lifecycle involves defining the objectives of the project and understanding the requirements, including compliance and regulatory considerations. This phase ensures the ML solution aligns with legal and organizational standards before proceeding to technical stages like data collection or model training.
Exact Extract from AWS AI Documents:
From the AWS AI Practitioner Learning Path:
" The business goal identification phase involves defining the problem to be solved, identifying success metrics, and determining compliance and regulatory requirements to ensure the ML solution adheres to legal and organizational standards. " (Source: AWS AI Practitioner Learning Path, Module on Machine Learning Lifecycle) Detailed Explanation:
Option A: Feature engineeringFeature engineering involves creating or selecting features for model training, which occurs after compliance requirements are identified. It does not address regulatory concerns.
Option B: Model trainingModel training focuses on building the ML model using data, not on determining compliance or regulatory requirements.
Option C: Data collectionData collection involves gathering data for training, but compliance and regulatory requirements (e.g., data privacy laws) are defined earlier in the business goal identification phase.
Option D: Business goal identificationThis is the correct answer. This phase ensures that compliance and regulatory requirements are considered at the outset, shaping the entire ML project.
References:

AWS AI Practitioner Learning Path: Module on Machine Learning Lifecycle Amazon SageMaker Developer Guide: ML Workflow (https://docs.aws.amazon.com/sagemaker/latest/dg/how-it-works-mlconcepts.html)
AWS Well-Architected Framework: Machine Learning Lens (https://docs.aws.amazon.com/wellarchitected/latest/machine-learning-lens/)

QUESTION NO: 22
F1 점수는 기초 모델(FM) 성과의 맥락에서 무엇을 측정합니까?
A. 모델 정확도와 재현율.
B. 응답을 생성하는 모델 속도.
C. 모델 운영에 드는 재정적 비용.
D. 모델 계산의 에너지 효율성.
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

QUESTION NO: 23
한 회사가 과거 데이터가 포함된 여러 데이터셋을 보유하고 있습니다. 이 회사는 머신러닝 기술을 사용하여 각 데이터셋을 처리하고자 합니다.
다음 목록에서 각 데이터 세트에 맞는 머신러닝 기술을 선택하세요. 각 머신러닝 기술은 한 번씩만 선택하거나 전혀 선택하지 않아도 됩니다. (세 가지를 선택하세요.) 컴퓨터 비전, 자연어 처리(NLP), 강화 학습, 시계열 예측

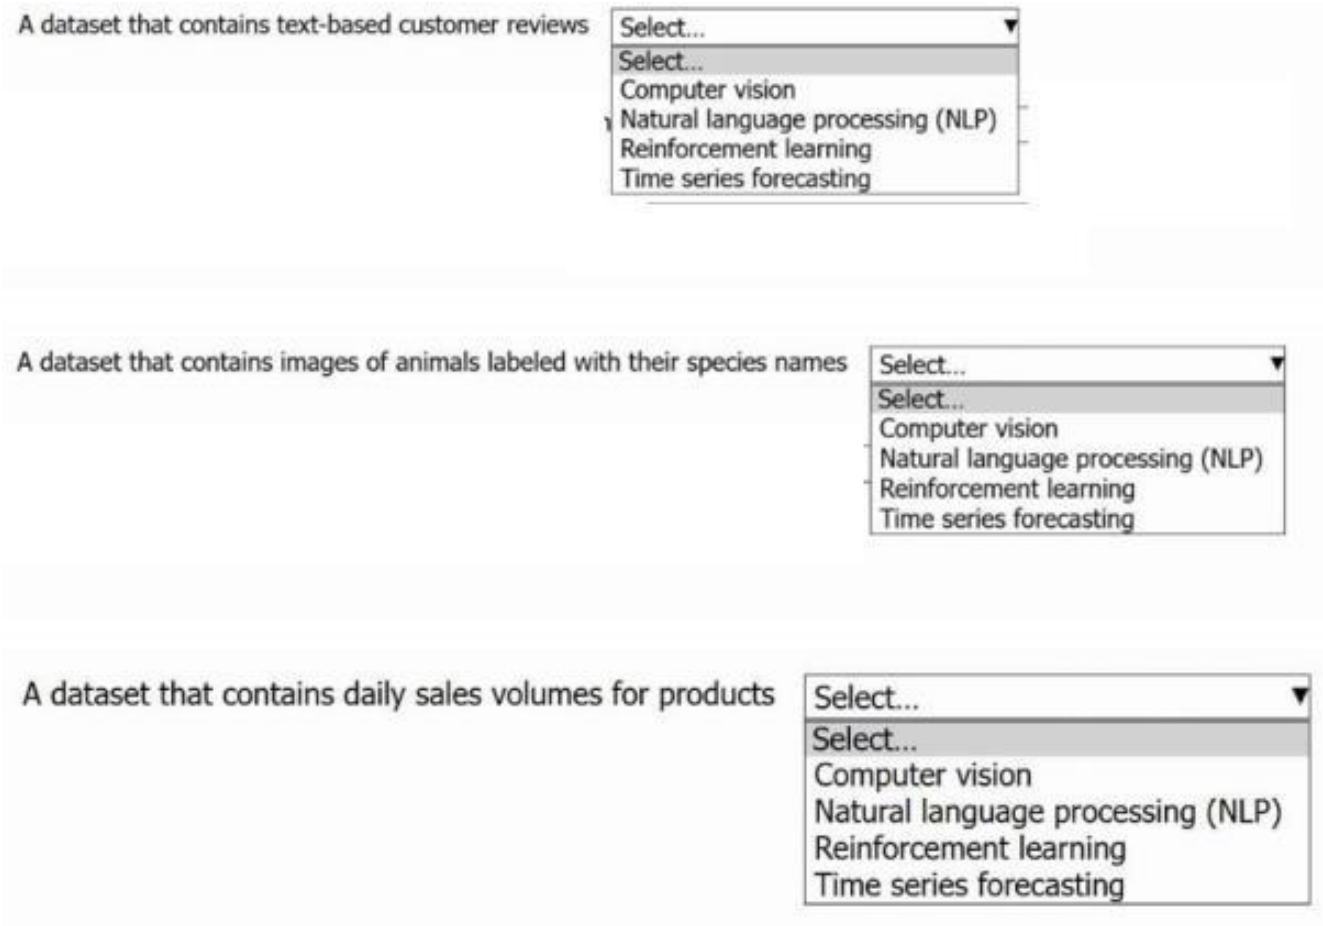
A dataset that contains text-based customer reviews
Select...
Select...
Computer vision
Natural language processing (NLP)
Reinforcement learning
Time series forecasting
A dataset that contains images of animals labeled with their species names
Select...
Select...
Computer vision
Natural language processing (NLP)
Reinforcement learning
Time series forecasting
A dataset that contains daily sales volumes for products
Select...
Select...
Computer vision
Natural language processing (NLP)
Reinforcement learning
Time series forecasting

Answer:

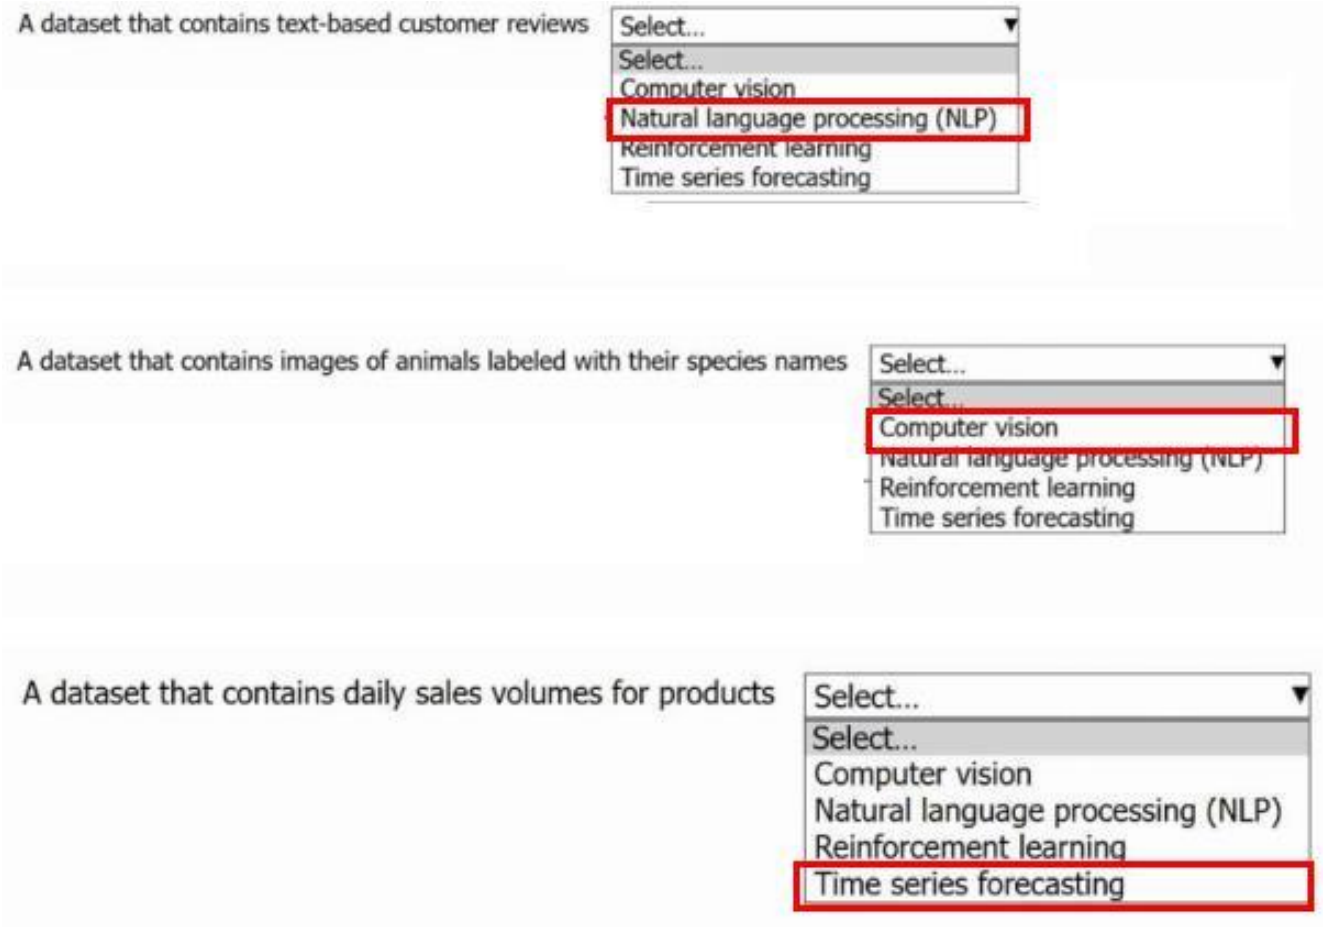
A dataset that contains text-based customer reviews
Select...
Select...
Computer vision
Natural language processing (NLP)
Reinforcement learning
Time series forecasting
A dataset that contains images of animals labeled with their species names
Select...
Select...
Computer vision
Natural language processing (NLP)
Reinforcement learning
Time series forecasting
A dataset that contains daily sales volumes for products
Select...
Select...
Computer vision
Natural language processing (NLP)
Reinforcement learning
Time series forecasting

Explanation:

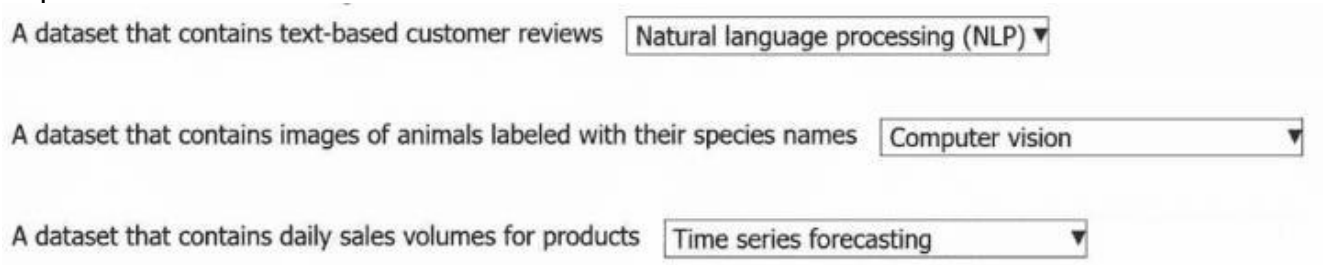

Dataset 1: A dataset that contains text-based customer reviews # Natural language processing (NLP) NLP is designed for analyzing text (sentiment analysis, text classification, etc.).
Dataset 2: A dataset that contains images of animals labeled with their species names # Computer vision Computer vision models classify or detect objects in images.
Dataset 3: A dataset that contains daily sales volumes for products # Time series forecasting Time series forecasting predicts future values based on historical sequential data (like sales, demand, stock prices).

QUESTION NO: 24
감정 분석은 어떤 AI 분야의 하위 집합인가요?
A. 컴퓨터 비전
B. 로봇공학
C. 자연어 처리(NLP)
D. 시계열 예측
Answer: C
Explanation:
Sentiment analysis is the task of determining the emotional tone or intent behind a body of text (positive, negative, neutral).
This falls under Natural Language Processing (NLP) because it deals with understanding and processing human language.
Computer vision relates to images, robotics to autonomous machines, and time series forecasting to predicting values from sequential data.
Reference:
AWS ML Glossary - NLP

QUESTION NO: 25
AI 실무자가 대규모 언어 모델(LLM)을 사용하여 마케팅 캠페인을 위한 콘텐츠를 만들고 있습니다. 생성된 콘텐츠는 그럴듯하고 사실처럼 들리지만 사실이 아닙니다.
LLM은 어떤 문제를 가지고 있나요?
A. 데이터 유출
B. 환각
C. 과적합
D. 과소적합
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 problemoften 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.

QUESTION NO: 26
한 회사가 기초 모델(FM)을 훈련하고 있습니다. 이 회사는 모델의 정확도를 특정 허용 수준까지 높이고 싶어합니다.
어떤 솔루션이 이러한 요구 사항을 충족시킬까요?
A. 배치 크기를 줄입니다.
B. 에포크를 늘립니다.
C. 에포크를 감소시킵니다.
D. 온도 매개변수를 증가시킵니다.
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.

QUESTION NO: 27
한 회사가 기초 모델(FM)을 사용하여 AI 모델을 개발하고 배포하려고 합니다.

어떤 AWS 서비스 또는 리소스가 최소한의 개발 노력으로 이러한 요구 사항을 충족할까요?
A. 아마존 기반암
B. 아마존 세이지메이커 AI
C. 아마존 베드락 파티락
D. 아마존 Q 개발자
Answer: A
Explanation:
Comprehensive and Detailed Explanation From Exact AWS AI documents:
Amazon Bedrock is the managed AWS service designed specifically to develop, customize, and deploy foundation models with minimal effort.
Amazon Bedrock:
Provides direct access to multiple foundation models
Eliminates infrastructure management
Supports customization, evaluation, and deployment
Why the other options are incorrect:
SageMaker AI (B) requires more setup and ML expertise.
PartyRock (C) is for experimentation, not deployment.
Amazon Q Developer (D) is a productivity assistant.
AWS AI document references:
Amazon Bedrock Overview
Foundation Model Deployment on AWS
Reducing Development Effort with Managed AI Services

QUESTION NO: 28
한 디자인 회사가 Amazon Bedrock의 기초 모델(FM)을 사용하여 다양한 프로젝트의 이미지를 생성하고 있습니다. 이 회사는 생성된 각 이미지가 얼마나 자세하거나 추상적으로 보이는지 제어하고 싶어합니다.
회사는 어떤 모델 매개변수를 수정해야 합니까?
A. 모델 체크포인트
B. 배치 크기
C. 생성 단계
D. 토큰 길이
Answer: C
Explanation:
The correct answer is C because in image generation tasks using foundation models like Stable Diffusion on Amazon Bedrock, the number of generation steps directly influences the fidelity and level of detail of the generated image. Fewer steps can produce more abstract or less defined images, while more steps allow the model to refine details, resulting in higher realism.
From AWS documentation:
" In diffusion-based image generation models, the number of inference steps (generation steps) determines how refined the final image is. Lower steps produce faster but less detailed outputs. Increasing steps results in more detailed and higher-quality images. " Explanation of other options:

A). Model checkpoint refers to saved versions of the model during training, not inference-time generation settings.
B). Batch size affects training/inference throughput, not image detail.
D). Token length is relevant for text models, not image generation.
Referenced AWS AI/ML Documents and Study Guides:
Amazon Bedrock Documentation - Model Parameters for Image Generation
Stable Diffusion on Amazon Bedrock - Developer Guide
AWS ML Specialty Guide - Generative AI Configuration

QUESTION NO: 29
한 회사가 고객 서비스 AI 어시스턴트를 위한 두 가지 기본 모델(FM)을 비교하고 있습니다. 회사는 유용성, 정확성 및 어조를 기준으로 FM을 평가하고자 합니다. 이를 위해 자동화되고 반복 가능하며 사람의 검토가 필요 없는 평가 기법을 찾고 있습니다.
어떤 평가 기법이 이러한 요구 사항을 충족할까요?
A. 문자열 일치
B. 요약 평가를 위한 회상 중심 보조 연구(ROUGE)
C. 판사로서의 LLM
D. 검색 증강 생성(RAG)
Answer: C
Explanation:
AWS documentation describes LLM-as-a-judge as an automated evaluation technique where a large language model is used to assess the outputs of another model based on qualitative criteria such as helpfulness, correctness, tone, and alignment with expectations . This approach enables scalable and repeatable evaluations without requiring human reviewers.
In this scenario, the company needs to compare two foundation models across subjective dimensions that are difficult to measure using traditional metrics. LLM-as-a-judge allows the evaluator model to score or rank responses using predefined evaluation prompts and criteria, ensuring consistent and automated assessment.
The other options do not meet the requirements. String matching and ROUGE focus on lexical similarity and are unsuitable for evaluating tone or helpfulness in customer service interactions. Retrieval Augmented Generation is an architectural pattern, not an evaluation technique.
AWS highlights LLM-as-a-judge as a practical approach for automated qualitative evaluation of generative AI outputs , making it the correct choice.

QUESTION NO: 30
한 금융 회사가 대출 신청 결과를 예측하는 생성형 AI 모델을 학습시키고 있습니다. 학습 데이터셋은 규모가 작습니다. 데이터셋은 대출 신청자를 "젊은층", "중년층", "노년층"으로 분류합니다. 데이터셋에 있는 대부분의 사람들은 "중년층"으로 분류됩니다. 이 회사는 학습 데이터셋에서 연령대 분류 특징을 제거합니다.
데이터셋 변경으로 인해 모델 동작에 어떤 변화가 나타날 가능성이 높습니까?
A. 이 모델은 젊은 연령층과 고령층의 결과를 부정확하게 예측할 것입니다.
B. 이 모델은 더 적은 훈련 데이터가 필요합니다.
C. 이 모델은 젊은 연령층에 대해서만 정확한 결과를 예측합니다.
D. 이 모델은 모든 연령대의 결과를 정확하게 예측합니다.

Answer: B

QUESTION NO: 31
한 회사가 대규모 언어 모델(LLM)을 사용하여 챗봇을 개발하려고 합니다. 이 챗봇은 고객의 제품 문의, 주문 추적 및 반품을 지원할 것입니다. 챗봇은 텍스트 입력과 이미지 입력을 처리하여 응답을 생성할 수 있어야 합니다.
어떤 AWS 서비스가 이러한 요구 사항을 충족합니까?
A. 아마존 기반암
B. 아마존 컴프리헨드
C. 아마존 Q
D. 아마존 레코그니션
Answer: A
Explanation:
Comprehensive and Detailed Explanation From Exact AWS AI documents:
Amazon Bedrock provides access to multimodal foundation models that can process both text and image inputs and generate intelligent responses.
Bedrock is designed for:
Building chatbots and conversational AI
Handling multimodal inputs
Integrating LLMs with enterprise applications
Why the other options are incorrect:
Amazon Comprehend (B) performs text analysis only.
Amazon Q (C) is a managed assistant, not a general chatbot platform.
Amazon Rekognition (D) analyzes images but does not generate conversational responses.
AWS AI document references:
Amazon Bedrock Overview
Multimodal Foundation Models on AWS
Building Chatbots with Amazon Bedrock

QUESTION NO: 32
한 회사가 실험 환경에서 생성적 AI 애플리케이션에 대해 알아보고자 합니다. 이 요구 사항을 가장 비용 효율적으로 충족할 수 있는 솔루션은 무엇일까요?
A. Amazon Q 개발자
B. Amazon SageMaker JumpStart
C. 아마존 베드락 파티락
D. 아마존 Q 비즈니스
Answer: C
Explanation:
The correct answer is Amazon Bedrock PartyRock, a playground for building and experimenting with generative AI apps in a low-cost, no-code environment. PartyRock is designed for innovation and learning. It enables users to try out prompts, LLM apps, and templates using Amazon Bedrock under a free-tier friendly setup. According to AWS, PartyRock abstracts infrastructure and allows rapid prototyping using models from Bedrock providers. This makes it ideal for early experimentation, especially for non-developers or

those not ready to invest in full production pipelines. In contrast, Amazon Q Developer is for software engineering tasks, SageMaker JumpStart focuses on deploying ML models, and Q Business targets enterprise knowledge workers. None of those are as cost-effective and experimental-focused as PartyRock.
Referenced AWS AI/ML Documents and Study Guides:
Amazon Bedrock Documentation - PartyRock Overview
AWS Generative AI Learning Path - Getting Started Tools

QUESTION NO: 33
한 회사가 인공지능(AI)을 이용해 고객의 관심사와 연령에 따라 장난감을 추천하는 웹사이트를 구축했습니다. 그런데 이 회사는 AI가 성별에 따른 고정관념에 기반한 장난감을 추천하는 경향이 있다는 것을 발견했습니다.
해당 기업이 편향성을 조사하기 위해 어떤 AWS 서비스 또는 기능을 사용해야 할까요?
A. 아마존 레코그니션
B. 아마존 Q 개발자
C. 아마존 컴프리헨드
D. 아마존 세이지메이커 클래리파이
Answer: D
Explanation:
Comprehensive and Detailed Explanation From Exact AWS AI documents:
Amazon SageMaker Clarify is designed to detect and explain bias in ML models and datasets.
AWS Responsible AI guidance recommends Clarify to:
Identify bias in predictions
Analyze feature attribution
Support fairness and ethical AI practices
Why the other options are incorrect:
Rekognition (A) analyzes images, not recommendation bias.
Amazon Q Developer (B) assists developers with code.
Comprehend (C) performs NLP tasks, not bias analysis.
AWS AI document references:
Amazon SageMaker Clarify Documentation
Detecting Bias in AI Systems
Responsible AI on AWS

QUESTION NO: 34
한 회사가 애플리케이션에 새로운 기능을 도입하려고 합니다. 이 기능은 출력 메시지의 스타일을 개선합니다. 회사는 이 기능을 구현하기 위해 Amazon Bedrock에서 대규모 언어 모델(LLM)을 미세 조정할 예정입니다. 이러한 요구 사항을 충족하려면 어떤 유형의 데이터가 필요할까요?
A. 입력 메시지만 샘플
B. 출력 메시지만 샘플
C. 입력 및 출력 메시지 쌍의 샘플
D. 입력 및 출력 메시지의 별도 샘플
Answer: C

Explanation:
Fine-tuning requires paired input-output examples to teach the model how to respond to inputs with desired styled outputs.
Single inputs (A) or outputs (B) are insufficient.
Separate, unpaired samples (D) don't establish the input-output mapping.
Reference:
AWS Documentation - Preparing data for fine-tuning FMs

QUESTION NO: 35
다큐멘터리 영화 제작자가 더 많은 시청자에게 다가가고 싶어합니다. 이 영화 제작자는 영화에 여러 언어의 자막과 음성 해설을 자동으로 추가하려고 합니다.
이러한 요구 사항을 충족하는 단계 조합은 무엇입니까? (두 가지를 선택하세요.)
A. Amazon Transcribe 및 Amazon Translate를 사용하여 다른 언어로 자막을 생성합니다.
B. Amazon Textract 및 Amazon Translate를 사용하여 다른 언어로 자막을 생성합니다.
C. Amazon Polly를 사용하여 다른 언어로 음성 해설을 생성합니다.
D. Amazon Translate를 사용하여 다른 언어로 음성 해설을 생성합니다.
E. Amazon Textract를 사용하여 다른 언어로 음성 해설을 생성합니다.
Answer: A C
Explanation:
The correct answers are A and C because:
Amazon Transcribe converts spoken dialogue from video into text (captions or subtitles).
Amazon Translate translates that transcribed text into other languages.
Amazon Polly can then convert the translated text into lifelike speech for voice-overs in different languages.
From AWS documentation:
" Amazon Transcribe is used to create accurate speech-to-text transcripts. You can feed this text into Amazon Translate to support multilingual subtitles. To generate audio output, Amazon Polly provides neural text-to- speech in multiple languages. " Explanation of other options:
B). Amazon Textract is used for extracting text from documents/images, not audio or video, and is not applicable here.
D). Amazon Translate does not generate speech - it only translates text.
E). Amazon Textract does not generate audio.
Referenced AWS AI/ML Documents and Study Guides:
Amazon Transcribe Developer Guide - Media Transcription
Amazon Translate Developer Guide - Multilingual Support
Amazon Polly Documentation - Text-to-Speech in Multiple Languages
AWS ML Specialty Guide - Multimedia AI Workflows