

## [2025-November-NewBraindump2go AIF-C01 VCE Exam Questions Free][Q121-Q155]

2025/November Latest Braindump2go AIF-C01 Exam Dumps with PDF and VCE Free Updated Today! Following are some new Braindump2go AIF-C01 Real Exam Questions!QUESTION 121An AI practitioner wants to predict the classification of flowers based on petal length, petal width, sepal length, and sepal width.Which algorithm meets these requirements?A. K-nearest neighbors (k-NN)B. K-meanC. Autoregressive Integrated Moving Average (ARIMA)D. Linear regressionAnswer: AQUESTION 122A company is using custom models in Amazon Bedrock for a generative AI application. The company wants to use a company managed encryption key to encrypt the model artifacts that the model customization jobs create.Which AWS service meets these requirements?A. AWS Key Management Service (AWS KMS)B. Amazon InspectorC. Amazon MacieD. AWS Secrets ManagerAnswer: AQUESTION 123A company wants to use large language models (LLMs) to produce code from natural language code comments.Which LLM feature meets these requirements?A. Text summarizationB. Text generationC. Text completionD. Text classificationAnswer: BQUESTION 124A company is introducing a mobile app that helps users learn foreign languages. The app makes text more coherent by calling a large language model (LLM). The company collected a diverse dataset of text and supplemented the dataset with examples of more readable versions. The company wants the LLM output to resemble the provided examples.Which metric should the company use to assess whether the LLM meets these requirements?A. Value of the loss functionB. Semantic robustnessC. Recall-Oriented Understudy for Gisting Evaluation (ROUGE) scoreD. Latency of the text generationAnswer: CQUESTION 125A company notices that its foundation model (FM) generates images that are unrelated to the prompts. The company wants to modify the prompt techniques to decrease unrelated images.Which solution meets these requirements?A. Use zero-shot prompts.B. Use negative prompts.C. Use positive prompts.D. Use ambiguous prompts.Answer: BQUESTION 126A company wants to use a large language model (LLM) to generate concise, feature-specific descriptions for the company's products.Which prompt engineering technique meets these requirements?A. Create one prompt that covers all products. Edit the responses to make the responses more specific, concise, and tailored to each product.B. Create prompts for each product category that highlight the key features. Include the desired output format and length for each prompt response.C. Include a diverse range of product features in each prompt to generate creative and unique descriptions.D. Provide detailed, product-specific prompts to ensure precise and customized descriptions.Answer: BQUESTION 127A company is developing an ML model to predict customer churn. The model performs well on the training dataset but does not accurately predict churn for new data.Which solution will resolve this issue?A. Decrease the regularization parameter to increase model complexity.B. Increase the regularization parameter to decrease model complexity.C. Add more features to the input data.D. Train the model for more epochs.Answer: BQUESTION 128A company is implementing intelligent agents to provide conversational search experiences for its customers. The company needs a database service that will support storage and queries of embeddings from a generative AI model as vectors in the database.Which AWS service will meet these requirements?A. Amazon AthenaB. Amazon Aurora PostgreSQLC. Amazon RedshiftD. Amazon EMRAnswer: BExplanation:The requirement is to identify an AWS database service that supports the storage and querying of embeddings (from a generative AI model) as vectors. Embeddings are typically high-dimensional numerical representations of data (e.g., text, images) used in AI applications like conversational search. The database must support vector storage and efficient vector similarity searches. Amazon Aurora PostgreSQL- Compatible Edition supports the pgvector extension, which enables efficient storage and similarity searches for vector embeddings. This makes it suitable for AI/ML workloads such as natural language processing and recommendation systems that rely on vector data.QUESTION 129A financial institution is building an AI solution to make loan approval decisions by using a foundation model (FM). For security and audit purposes, the company needs the AI solution's decisions to be explainable.Which factor relates to the explainability of the AI solution's decisions?A. Model complexityB. Training timeC. Number of hyperparametersD. Deployment timeAnswer: AExplanation:The financial institution needs an AI solution for loan approval decisions to be explainable for security and audit purposes. Explainability refers to the ability to understand and interpret how a model makes decisions. Model complexity directly impacts explainability: simpler models (e.g., logistic regression) are more interpretable, while complex models (e.g., deep neural networks) are harder to explain, often behaving like "black boxes".Model complexity affects the explainability of AI solutions. Simpler models, such as linear regression, are inherently more interpretable, while complex models, such as deep neural networks, may require additional tools like SageMaker Clarify to provide insights into their decision-making processes.QUESTION 130A company wants to use Amazon Bedrock. The company needs to review which security aspects the company is responsible for when using Amazon Bedrock.Which security aspect will the company be responsible for?A. Patching and updating the versions of Amazon BedrockB. Protecting the infrastructure that hosts Amazon BedrockC. Securing the

company's data in transit and at restD. Provisioning Amazon Bedrock within the company networkAnswer: CQUESTION 131A company wants to build a lead prioritization application for its employees to contact potential customers. The application must give employees the ability to view and adjust the weights assigned to different variables in the model based on domain knowledge and expertise.Which ML model type meets these requirements?A. Logistic regression modelB. Deep learning model built on principal componentsC. K-nearest neighbors (k-NN) modelD. Neural networkAnswer: AExplanation: The company needs an ML model for a lead prioritization application where employees can view and adjust the weights assigned to different variables based on domain knowledge. Logistic regression is a linear model that assigns interpretable weights to input features, making it easy for users to understand and modify these weights. This interpretability and adjustability make it suitable for the requirements.Logistic regression is a supervised learning algorithm used for classification tasks. It is highly interpretable, as it assigns weights to each feature, allowing users to understand and adjust the importance of different variables based on domain expertise.QUESTION 132 Which strategy will determine if a foundation model (FM) effectively meets business objectives?A. Evaluate the model's performance on benchmark datasets.B. Analyze the model's architecture and hyperparameters.C. Assess the model's alignment with specific use cases.D. Measure the computational resources required for model deployment. Answer: CQUESTION 133A social media company wants to use a large language model (LLM) to summarize messages. The company has chosen a few LLMs that are available on Amazon SageMaker JumpStart. The company wants to compare the generated output toxicity of these models. Which strategy gives the company the ability to evaluate the LLMs with the LEAST operational overhead?A. Crowd-sourced evaluationB. Automatic model evaluationC. Model evaluation with human workersD. Reinforcement learning from human feedback (RLHF)Answer: BQUESTION 134Which phase of the ML lifecycle determines compliance and regulatory requirements?A. Feature engineeringB. Model trainingC. Data collectionD. Business goal identificationAnswer: DExplanation: 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. 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.QUESTION 135A food service company wants to develop an ML model to help decrease daily food waste and increase sales revenue. The company needs to continuously improve the model's accuracy.Which solution meets these requirements?A. Use Amazon SageMaker and iterate with newer data.B. Use Amazon Personalize and iterate with historical data.C. Use Amazon CloudWatch to analyze customer orders.D. Use Amazon Rekognition to optimize the model. Answer: AQUESTION 136A company has developed an ML model to predict real estate sale prices. The company wants to deploy the model to make predictions without managing servers or infrastructure. Which solution meets these requirements?A. Deploy the model on an Amazon EC2 instance.B. Deploy the model on an Amazon Elastic Kubernetes Service (Amazon EKS) cluster.C. Deploy the model by using Amazon CloudFront with an Amazon S3 integration.D. Deploy the model by using an Amazon SageMaker endpoint. Answer: DQUESTION 137A company wants to develop an AI application to help its employees check open customer claims, identify details for a specific claim, and access documents for a claim.Which solution meets these requirements?A. Use Agents for Amazon Bedrock with Amazon Fraud Detector to build the application.B. Use Agents for Amazon Bedrock with Amazon Bedrock knowledge bases to build the application.C. Use Amazon Personalize with Amazon Bedrock knowledge bases to build the application.D. Use Amazon SageMaker to build the application by training a new ML model. Answer: BExplanation: The company wants an AI application to help employees check open customer claims, identify claim details, and access related documents. Agents for Amazon Bedrock can automate tasks by interacting with external systems, while Amazon Bedrock knowledge bases provide a repository of information (e.g., claim details and documents) that the agent can query to respond to employee requests, making this the best solution. Agents for Amazon Bedrock enable developers to build applications that can perform tasks by interacting with external systems and data sources. When paired with Amazon Bedrock knowledge bases, agents can access structured and unstructured data, such as documents or databases, to provide detailed responses for use cases like customer service or claims management.QUESTION 138A manufacturing company uses AI to inspect products and find any damages or defects.Which type of AI application is the company using?A. Recommendation systemB. Natural language processing (NLP)C. Computer visionD. Image processingAnswer: CExplanation: The manufacturing company uses AI to inspect products for damages or defects, which involves analyzing visual data (e.g., images or videos of products). This task falls under computer vision, a type of AI application that enables machines to interpret and understand visual information, such as identifying defects in manufacturing. Computer vision enables machines to interpret and understand visual data from the world, such as images or videos. Common applications include defect detection in manufacturing, where AI models analyze product images to identify damages or anomalies.QUESTION 139A company wants to create an ML

model to predict customer satisfaction. The company needs fully automated model tuning. Which AWS service meets these requirements?A. Amazon PersonalizeB. Amazon SageMakerC. Amazon AthenaD. Amazon ComprehendAnswer: B  
QUESTION 140 Which technique can a company use to lower bias and toxicity in generative AI applications during the post-processing ML lifecycle?A. Human-in-the-loopB. Data augmentationC. Feature engineeringD. Adversarial trainingAnswer: A  
QUESTION 141 A bank has fine-tuned a large language model (LLM) to expedite the loan approval process. During an external audit of the model, the company discovered that the model was approving loans at a faster pace for a specific demographic than for other demographics. How should the bank fix this issue MOST cost-effectively?A. Include more diverse training data. Fine-tune the model again by using the new data.B. Use Retrieval Augmented Generation (RAG) with the fine-tuned model.C. Use AWS Trusted Advisor checks to eliminate bias.D. Pre-train a new LLM with more diverse training data. Answer: A  
QUESTION 142 A company needs to log all requests made to its Amazon Bedrock API. The company must retain the logs securely for 5 years at the lowest possible cost. Which combination of AWS service and storage class meets these requirements? (Choose two.)A. AWS CloudTrailB. Amazon CloudWatchC. AWS Audit ManagerD. Amazon S3 Intelligent-TieringE. Amazon S3 StandardAnswer: AD  
QUESTION 143 An ecommerce company wants to improve search engine recommendations by customizing the results for each user of the company's ecommerce platform. Which AWS service meets these requirements?A. Amazon PersonalizeB. Amazon KendraC. Amazon RekognitionD. Amazon TranscribeAnswer: A  
Explanation: The ecommerce company wants to improve search engine recommendations by customizing results for each user. Amazon Personalize is a machine learning service that enables personalized recommendations, tailoring search results or product suggestions based on individual user behavior and preferences, making it the best fit for this requirement. Amazon Personalize enables developers to build applications with personalized recommendations, such as customized search results or product suggestions, by analyzing user behavior and preferences to deliver tailored experiences.  
QUESTION 144 A hospital is developing an AI system to assist doctors in diagnosing diseases based on patient records and medical images. To comply with regulations, the sensitive patient data must not leave the country the data is located in. Which data governance strategy will ensure compliance and protect patient privacy?A. Data residencyB. Data qualityC. Data discoverabilityD. Data enrichmentAnswer: A  
QUESTION 145 A company needs to monitor the performance of its ML systems by using a highly scalable AWS service. Which AWS service meets these requirements?A. Amazon CloudWatchB. AWS CloudTrailC. AWS Trusted AdvisorD. AWS ConfigAnswer: A  
QUESTION 146 An AI practitioner is developing a prompt for an Amazon Titan model. The model is hosted on Amazon Bedrock. The AI practitioner is using the model to solve numerical reasoning challenges. The AI practitioner adds the following phrase to the end of the prompt: "Ask the model to show its work by explaining its reasoning step by step." Which prompt engineering technique is the AI practitioner using?A. Chain-of-thought promptingB. Prompt injectionC. Few-shot promptingD. Prompt templatingAnswer: A  
QUESTION 147 Which AWS service makes foundation models (FMs) available to help users build and scale generative AI applications?A. Amazon Q DeveloperB. Amazon BedrockC. Amazon KendraD. Amazon ComprehendAnswer: B  
Explanation: Amazon Bedrock is a fully managed service that provides access to foundation models (FMs) from various providers, enabling users to build and scale generative AI applications. It simplifies the process of integrating FMs into applications for tasks like text generation, chatbots, and more. Amazon Bedrock is a fully managed service that makes foundation models (FMs) from leading AI providers available through a single API, enabling developers to build and scale generative AI applications with ease.  
QUESTION 148 A company is building a mobile app for users who have a visual impairment. The app must be able to hear what users say and provide voice responses. Which solution will meet these requirements?A. Use a deep learning neural network to perform speech recognition.B. Build ML models to search for patterns in numeric data.C. Use generative AI summarization to generate human-like text.D. Build custom models for image classification and recognition. Answer: A  
Explanation: The mobile app for users with visual impairment needs to hear user speech and provide voice responses, requiring speech-to-text (speech recognition) and text-to-speech capabilities. Deep learning neural networks are widely used for speech recognition tasks, as they can effectively process and transcribe spoken language. AWS services like Amazon Transcribe, which uses deep learning for speech recognition, can fulfill this requirement by converting user speech to text, and Amazon Polly can generate voice responses. Amazon Transcribe uses deep learning neural networks to perform automatic speech recognition (ASR), converting spoken language into text with high accuracy. This is ideal for applications requiring voice input, such as accessibility features for visually impaired users.  
QUESTION 149 A company wants to enhance response quality for a large language model (LLM) for complex problem-solving tasks. The tasks require detailed reasoning and a step-by-step explanation process. Which prompt engineering technique meets these requirements?A. Few-shot promptingB. Zero-shot promptingC. Directional stimulus promptingD. Chain-of-thought promptingAnswer: D  
Explanation: The company wants to enhance the response quality of an LLM for complex problem-solving tasks requiring detailed reasoning and step-by-step explanations. Chain-of-thought prompting encourages the LLM

to break down the problem into intermediate steps, providing a clear reasoning process before arriving at the final answer, which is ideal for this requirement. Chain-of-thought prompting improves the reasoning capabilities of large language models by encouraging them to break down complex tasks into intermediate steps, providing a step-by-step explanation that leads to the final answer. This technique is particularly effective for problem-solving tasks requiring detailed reasoning.

**QUESTION 150**A company wants to keep its foundation model (FM) relevant by using the most recent data. The company wants to implement a model training strategy that includes regular updates to the FM. Which solution meets these requirements?

- A. Batch learning
- B. Continuous pre-training
- C. Static training
- D. Latent training

**Answer: B**QUESTION 151Which option is a characteristic of AI governance frameworks for building trust and deploying human-centered AI technologies?

- A. Expanding initiatives across business units to create long-term business value
- B. Ensuring alignment with business standards, revenue goals, and stakeholder expectations
- C. Overcoming challenges to drive business transformation and growth
- D. Developing policies and guidelines for data, transparency, responsible AI, and compliance

**Answer: D**Explanation: AI governance frameworks aim to build trust and deploy human-centered AI technologies by establishing guidelines and policies for data usage, transparency, responsible AI practices, and compliance with regulations. This ensures ethical and accountable AI development and deployment. AI governance frameworks establish trust in AI technologies by developing policies and guidelines for data management, transparency, responsible AI practices, and compliance with regulatory requirements, ensuring human-centered and ethical AI deployment.

**QUESTION 152**An ecommerce company is using a generative AI chatbot to respond to customer inquiries. The company wants to measure the financial effect of the chatbot on the company's operations. Which metric should the company use?

- A. Number of customer inquiries handled
- B. Cost of training AI models
- C. Cost for each customer conversation
- D. Average handled time (AHT)

**Answer: C**QUESTION 153A company wants to find groups for its customers based on the customers' demographics and buying patterns. Which algorithm should the company use to meet this requirement?

- A. K-nearest neighbors (k-NN)
- B. K-means
- C. Decision tree
- D. Support vector machine

**Answer: B**QUESTION 154A company's large language model (LLM) is experiencing hallucinations. How can the company decrease hallucinations?

- A. Set up Agents for Amazon Bedrock to supervise the model training
- B. Use data pre-processing and remove any data that causes hallucinations
- C. Decrease the temperature inference parameter for the model
- D. Use a foundation model (FM) that is trained to not hallucinate

**Answer: C**Explanation: Hallucinations in large language models (LLMs) occur when the model generates outputs that are factually incorrect, irrelevant, or not grounded in the input data. To mitigate hallucinations, adjusting the model's inference parameters, particularly the temperature, is a well-documented approach in AWS AI Practitioner resources. The temperature parameter controls the randomness of the model's output. A lower temperature makes the model more deterministic, reducing the likelihood of generating creative but incorrect responses, which are often the cause of hallucinations. The temperature parameter controls the randomness of the generated text. Higher values (e.g., 0.8 or above) increase creativity but may lead to less coherent or factually incorrect outputs, while lower values (e.g., 0.2 or 0.3) make the output more focused and deterministic, reducing the likelihood of hallucinations.

**QUESTION 155**A company is using a large language model (LLM) on Amazon Bedrock to build a chatbot. The chatbot processes customer support requests. To resolve a request, the customer and the chatbot must interact a few times. Which solution gives the LLM the ability to use content from previous customer messages?

- A. Turn on model invocation logging to collect messages
- B. Add messages to the model prompt
- C. Use Amazon Personalize to save conversation history
- D. Use Provisioned Throughput for the LLM

**Answer: B**Explanation: The company is building a chatbot using an LLM on Amazon Bedrock, and the chatbot needs to use content from previous customer messages to resolve requests. Adding previous messages to the model prompt (also known as providing conversation history) enables the LLM to maintain context across interactions, allowing it to respond coherently based on the ongoing conversation.

To enable a large language model (LLM) to maintain context in a conversation, you can include previous messages in the model prompt. This approach, often referred to as providing conversation history, allows the LLM to generate responses that are contextually relevant to prior interactions.

Resources From: [1.2025 Latest Braindump2go AIF-C01 Exam Dumps \(PDF & VCE\) Free Share:](https://www.braindump2go.com/aif-c01.html)

<https://www.braindump2go.com/aif-c01.html> 2.2025 Latest Braindump2go AIF-C01 PDF and AIF-C01 VCE Dumps Free Share:

<https://drive.google.com/drive/folders/1DWDPeocDJXhrT9CuX60G8X4OR-ZT7WE8?usp=sharing> 3.2025 Free

**Braindump2go AIF-C01 Exam Questions Download:**

[https://www.braindump2go.com/free-online-pdf/AIF-C01-VCE-Dumps\(121-155\).pdf](https://www.braindump2go.com/free-online-pdf/AIF-C01-VCE-Dumps(121-155).pdf) Free Resources from

**Braindump2go, We Devoted to Helping You 100% Pass All Exams!**