

Microsoft

Exam Questions AI-102

Designing and Implementing an Azure AI Solution



NEW QUESTION 1

HOTSPOT - (Topic 3)

You are creating an enrichment pipeline that will use Azure Cognitive Search. The knowledge store contains unstructured JSON data and scanned PDF documents that contain text.

Which projection type should you use for each data type? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

JSON data:	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>File projection</p> <p>Object projection</p> <p>Table projection</p> </div> </div>
Scanned data:	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <div style="padding: 2px;"> <p>File projection</p> <p>Object projection</p> <p>Table projection</p> </div> </div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Object projection

Object projections are JSON representations of the enrichment tree that can be sourced from any node.

Box 2: File projection

File projections are similar to object projections and only act on the normalized_images collection.

NEW QUESTION 2

- (Topic 3)

You are building a chatbot.

You need to configure the bot to guide users through a product setup process. Which type of dialog should you use?

- A. component
- B. waterfall
- C. adaptive
- D. action

Answer: B

NEW QUESTION 3

- (Topic 3)

You have a Video Indexer service that is used to provide a search interface over company videos on your company's website.

You need to be able to search for videos based on who is present in the video. What should you do?

- A. Create a person model and associate the model to the videos.
- B. Create person objects and provide face images for each object.
- C. Invite the entire staff of the company to Video Indexer.
- D. Edit the faces in the videos.
- E. Upload names to a language model.

Answer: A

Explanation:

Video Indexer supports multiple Person models per account. Once a model is created, you can use it by providing the model ID of a specific Person model when uploading/indexing or reindexing a video. Training a new face for a video updates the specific custom model that the video was associated with.

Note: Video Indexer supports face detection and celebrity recognition for video content. The celebrity recognition feature covers about one million faces based on commonly requested data source such as IMDB, Wikipedia, and top LinkedIn influencers. Faces that aren't recognized by the celebrity recognition feature are detected but left unnamed. Once you label a face with a name, the face and name get added to your account's Person model. Video Indexer will then recognize this face in your future videos and past videos.

Reference:

<https://docs.microsoft.com/en-us/azure/media-services/video-indexer/customize-person-model-with-api>

NEW QUESTION 4

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

The massively parallel processing (MPP) engine of Azure Synapse Analytics

- distributes processing across compute nodes.
- distributes processing across control nodes.
- redirects client connections across compute nodes.
- redirects client connections across control nodes.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

The massively parallel processing (MPP) engine of Azure Synapse Analytics

- distributes processing across compute nodes.
- distributes processing across control nodes.
- redirects client connections across compute nodes.
- redirects client connections across control nodes.

NEW QUESTION 5

- (Topic 3)

You are building a chatbot for a travel agent. The chatbot will use the Azure OpenAI GPT 3.5 model and will be used to make travel reservations. You need to maximize the accuracy of the responses from the chatbot. What should you do?

- A. Configure the model to include data from the travel agent's database.
- B. Set the Top P parameter for the model to 0.
- C. Set the Temperature parameter for the model to 0.
- D. Modify the system message used by the model to specify that the answers must be accurate.

Answer: A

NEW QUESTION 6

- (Topic 3)

What should you use to automatically delete blobs from Azure Blob Storage?

- A. the change feed
- B. a lifecycle management policy
- C. soft delete
- D. archive storage

Answer: D

NEW QUESTION 7

- (Topic 3)

You have an Azure subscription that contains an Anomaly Detector resource. You deploy a Docker host server named Server 1 to the on-premises network. You need to host an instance of the Anomaly Detector service on Server 1. Which parameter should you include in the docker run command?

- A. Fluentd
- B. Billing
- C. Http Proxy
- D. Mounts

Answer: B

NEW QUESTION 8

HOTSPOT - (Topic 3)

You have a bot that was built by using the Microsoft Bot Framework composer as shown in the following exhibit.



Use the drop-down menus 10 select the answer choice that completes each statement based on the information presented in the graphic.
 NOTE: Each correct selection is worth one point.

Answer Area

If a user asks "what is the weather like in New York", the bot will [answer choice].

- identify New York as a city entity
- change to a different dialog
- identify New York as a city entity
- identify New York as a state entity
- respond with the weather in Seattle

The GetWeather dialog uses a [answer choice] trigger.

- Language Understanding Intent recogn...
- Custom events
- Dialog events
- Language Understanding Intent recognized
- QnA Intent recognized

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

If a user asks "what is the weather like in New York", the bot will [answer choice].

- identify New York as a city entity
- change to a different dialog
- identify New York as a city entity
- identify New York as a state entity
- respond with the weather in Seattle

The GetWeather dialog uses a [answer choice] trigger.

- Language Understanding Intent recogn...
- Custom events
- Dialog events
- Language Understanding Intent recognized
- QnA Intent recognized

NEW QUESTION 9

- (Topic 3)

You have an Azure subscription that contains a Language service resource named ta1 and a virtual network named vnet1. You need to ensure that only resources in vnet1 can access ta1. What should you configure?

- A. a network security group (NSG) for vnet1
- B. Azure Firewall for vnet1
- C. the virtual network settings for ta 1
- D. a Language service container for ta1

Answer: C

NEW QUESTION 10

- (Topic 3)

You have a Language Understanding resource named lu1. You build and deploy an Azure bot named bot1 that uses lu1. You need to ensure that bot1 adheres to the Microsoft responsible AI principle of inclusiveness. How should you extend bot1?

- A. Implement authentication for bot1.
- B. Enable active learning for lu1.
- C. Host lu1 in a container.
- D. Add Direct Line Speech to bot1.

Answer: D

Explanation:

Inclusiveness: AI systems should empower everyone and engage people.

Direct Line Speech is a robust, end-to-end solution for creating a flexible, extensible voice assistant. It is powered by the Bot Framework and its Direct Line Speech channel, that is optimized for voice-in, voice-out interaction with bots.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/speech-service/direct-line-speech>

NEW QUESTION 10

- (Topic 3)

You are developing a monitoring system that will analyze engine sensor data, such as rotation speed, angle, temperature, and pressure. The system must generate an alert in response to atypical values. What should you include in the solution?

- A. Application Insights in Azure Monitor
- B. metric alerts in Azure Monitor
- C. Multivariate Anomaly Detection
- D. Univariate Anomaly Detection

Answer: C

NEW QUESTION 14

- (Topic 3)

You have a 20-GB file named file1, avi that is stored on a local drive. You need to index file1, avi by using the Azure Video indexer website. What should you do first?

- A. Upload File1 avi to an Azure storage queue.
- B. upload File1.avi to the www.youtube.com seepage
- C. Upload file1 avi to the Azure video indexer website.
- D. Upload file1 avi to Microsoft OneDrive.

Answer: D

Explanation:

This is because the Azure Video Indexer website allows you to upload videos from a URL or from your file system, but there are some limitations and considerations for each option1.

If you upload from your file system, the size of the file is limited to 2 GB, which is less than the size of file1.avi (20 GB). Therefore, this option is not feasible.

If you upload from a URL, the size of the file is limited to 30 GB, which is enough for file1.avi. However, the URL must be publicly accessible and valid, and the file must be accessible. You cannot use URLs from streaming services such as YouTube1. Therefore, options A and B are not valid.

The best option is to upload file1.avi to a cloud storage service such as Microsoft OneDrive, and then generate a shareable link for the file. You can then paste the link in the

Azure Video Indexer website and upload the video from the URL. This way, you can avoid the file size limitation and ensure that the file is accessible and valid2.

NEW QUESTION 19

HOTSPOT - (Topic 3)

You have an Azure subscription that has the following configurations:

- Subscription ID: 8d3591aa-96b8-4737-ad09-00f9b1ed35ad
- Tenant ID: 3edfe572-cb54-3ced-ae12-c5c177f39a12


You plan to create a resource that will perform sentiment analysis and optical character recognition (OCR).

You need to use an HTTP request to create the resource in the subscription. The solution must use a single key and endpoint.

How should you complete the request? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area




- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 20

- (Topic 3)

You are building a Chatbot by using the Microsoft Bot Framework SDK. The bot will be used to accept food orders from customers and allow the customers to

customize each food item. You need to configure the bot to ask the user for additional input based on the type of item ordered. The solution must minimize development effort. Which two types of dialogs should you use? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. adaptive
- B. action
- C. waterfall
- D. prompt
- E. input

Answer: BC

NEW QUESTION 25

- (Topic 3)

You need to build a chatbot that meets the following requirements:

- ? Supports chit-chat, knowledge base, and multilingual models
- ? Performs sentiment analysis on user messages
- ? Selects the best language model automatically

What should you integrate into the chatbot?

- A. QnA Maker, Language Understanding, and Dispatch
- B. Translator, Speech, and Dispatch
- C. Language Understanding, Text Analytics, and QnA Maker
- D. Text Analytics, Translator, and Dispatch

Answer: C

Explanation:

Language Understanding: An AI service that allows users to interact with your applications, bots, and IoT devices by using natural language.

QnA Maker is a cloud-based Natural Language Processing (NLP) service that allows you to create a natural conversational layer over your data. It is used to find the most appropriate answer for any input from your custom knowledge base (KB) of information.

Text Analytics: Mine insights in unstructured text using natural language processing (NLP)—no machine learning expertise required. Gain a deeper understanding of customer opinions with sentiment analysis. The Language Detection feature of the Azure Text Analytics REST API evaluates text input

Reference:

<https://azure.microsoft.com/en-us/services/cognitive-services/text-analytics/> <https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/overview/overview>

NEW QUESTION 29

HOTSPOT - (Topic 3)

You are developing a service that records lectures given in English (United Kingdom).

You have a method named AppendToTranscriptFile that takes translated text and a language identifier.

You need to develop code that will provide transcripts of the lectures to attendees in their respective language. The supported languages are English, French, Spanish, and German.

How should you complete the code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

```
static async Task TranslateSpeechAsync()
{
    var config = SpeechTranslationConfig.FromSubscription("69cad5cc-0ab3-4704-bdff-afbf4aa07d85", "uksouth");

    var lang = new List<string>
    {
        {"en-GB"},
        {"fr", "de", "es"},
        {"French", "Spanish", "German"}
    };

    config.SpeechRecognitionLanguage = "en-GB";
    lang.ForEach(config.AddTargetLanguage);

    using var audioConfig = AudioConfig.FromDefaultMicrophoneInput();
    using var recognizer = new TranslationRecognizer(config, audioConfig);

    var result = await recognizer.RecognizeOnceAsync();
    if (result.Reason == ResultReason.TranslatedSpeech)
    {

```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: {"fr", "de", "es"}

A common task of speech translation is to specify target translation languages, at least one is required but multiples are supported. The following code snippet sets both French and German as translation language targets.

static async Task TranslateSpeechAsync()

```
{
    var translationConfig =
    SpeechTranslationConfig.FromSubscription(SPEECH SUBSCRIPTION KEY, SPEECH SERVICE REGION);
    translationConfig.SpeechRecognitionLanguage = "it-IT";
    // Translate to languages. See, https://aka.ms/speech/sttt-languages translationConfig.AddTargetLanguage("fr"); translationConfig.AddTargetLanguage("de");

```

```

}
Box 2: TranslationRecognizer
After you've created a SpeechTranslationConfig, the next step is to initialize a TranslationRecognizer.
Example code:
static async Task TranslateSpeechAsync()
{
var translationConfig =
SpeechTranslationConfig.FromSubscription(SPEECH SUBSCRIPTION KEY, SPEECH SERVICE REGION);
var fromLanguage = "en-US";
var toLanguages = new List<string> { "it", "fr", "de" }; translationConfig.SpeechRecognitionLanguage = fromLanguage;
toLanguages.ForEach(translationConfig.AddTargetLanguage);
using var recognizer = new TranslationRecognizer(translationConfig);
}

```

NEW QUESTION 31

- (Topic 3)
 What should you use to build a Microsoft Power Bi paginated report?

- A. Power BI Report Builder
- B. Charcicator
- C. Power BI Desktop
- D. the Power BI service

Answer: A

NEW QUESTION 32

- (Topic 3)
 Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You have a chatbot that uses question answering in Azure Cognitive Service for Language Users report that the responses of the chatbot lack formality when answering spurious questions. You need to ensure that the chatbot provides formal responses to spurious questions. Solution: From Language Studio, you change the chitchat source to qna_chitchit_friindly.tsv. and then retrain and republish the model. Does this meet the goal?



- A. Yes
- B. No

Answer: B

NEW QUESTION 33

HOTSPOT - (Topic 3)
 You are designing a conversation flow to be used in a chatbot. You need to test the conversation flow by using the Microsoft Bot Framework Emulator. How should you complete the .chat file? To answer, select the appropriate options in the answer area.
 NOTE: Each correct selection is worth one point.

```

user=User1
bot=watchbot
user: I want a new watch.
bot: [  ][Delay=3000]
bot: I can help you with that! Let me see what I can find.
bot: Here's what I found.
bot:
[AttachmentLayout=  ]
[Attachment=https://contoso.blob.core.windows.net/watch01.jpg]
[Attachment=https://contoso.blob.core.windows.net/watch02.jpg]
user: I like the first one.
bot: Sure, pulling up more information.
bot: [Attachment=cards\watchProfileCard.json  ]
user: That's nice! Thank you.
bot: Sure, you are most welcome!

```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

```
user=User1
bot=watchbot
user: I want a new watch.
bot: [ Attachment
      ConversationUpdate
      Typing ] [Delay=3000]
```

bot: I can help you with that! Let me see what I can find.
 bot: Here's what I found.
 bot:

```
[AttachmentLayout= adaptivecard
                  carousel
                  thumbnail ]
```

```
[Attachment=https://contoso.blob.core.windows.net/watch01.jpg]
[Attachment=https://contoso.blob.core.windows.net/watch02.jpg]
```

user: I like the first one.
 bot: Sure, pulling up more information.

```
bot: [Attachment=cards\watchProfileCard.json
      adaptivecard
      carousel
      list ]
```

user: That's nice! Thank you.
 bot: Sure, you are most welcome!

NEW QUESTION 35

- (Topic 3)

You have an Azure Cognitive Search solution and a collection of handwritten letters stored as JPEG files. You plan to index the collection. The solution must ensure that queries can be performed on the contents of the letters. You need to create an indexer that has a skillset. Which skill should you include?

- A. key phrase extraction
- B. optical character recognition (OCR)
- C. document extraction
- D. image analysis

Answer: B

NEW QUESTION 39

- (Topic 3)

Your company wants to reduce how long it takes for employees to log receipts in expense reports. All the receipts are in English. You need to extract top-level information from the receipts, such as the vendor and the transaction total. The solution must minimize development effort. Which Azure Cognitive Services service should you use?

- A. Custom Vision
- B. Personalizer
- C. Form Recognizer
- D. Computer Vision

Answer: B

Explanation:

Azure Form Recognizer is a cognitive service that lets you build automated data processing software using machine learning technology. Identify and extract text, key/value pairs, selection marks, tables, and structure from your documents—the service outputs structured data that includes the relationships in the original file, bounding boxes, confidence and more.

Form Recognizer is composed of custom document processing models, prebuilt models for invoices, receipts, IDs and business cards, and the layout model.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/form-recognizer>

NEW QUESTION 41

- (Topic 3)

You are building a bot on a local computer by using the Microsoft Bot Framework. The bot will use an existing Language Understanding model. You need to translate the Language Understanding model locally by using the Bot Framework CLI. What should you do first?

- A. From the Language Understanding portal, clone the model.
- B. Export the model as an .lu file.
- C. Create a new Speech service.
- D. Create a new Language Understanding service.

Answer: B

Explanation:

You might want to manage the translation and localization for the language understanding content for your bot independently. Translate command in the @microsoft/bf-lu library takes advantage of the Microsoft text translation API to automatically machine translate .lu files to one or more than 60+ languages supported by the Microsoft text translation cognitive service. What is translated? An .lu file and optionally translate Comments in the lu file LU reference link texts List of .lu files under a specific path. Reference: <https://github.com/microsoft/botframework-cli/blob/main/packages/luis/docs/translate-command.md>

NEW QUESTION 45

- (Topic 3)
 Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You create a web app named app1 that runs on an Azure virtual machine named vm1. Vm1 is on an Azure virtual network named vnet1. You plan to create a new Azure Cognitive Search service named service1. You need to ensure that app1 can connect directly to service1 without routing traffic over the public internet. Solution: You deploy service1 and a private endpoint to vnet1. Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

A private endpoint is a network interface that uses a private IP address from your virtual network. This network interface connects you privately and securely to a service powered by Azure Private Link. By enabling a private endpoint, you're bringing the service into your virtual network. The service could be an Azure service such as:
 ? Azure Storage
 ? Azure Cosmos DB
 ? Azure SQL Database
 ? Your own service using a Private Link Service. Reference: <https://docs.microsoft.com/en-us/azure/private-link/private-endpoint-overview>

NEW QUESTION 47

DRAG DROP - (Topic 3)
 You are developing a call to the Face API. The call must find similar faces from an existing list named employeefaces. The employeefaces list contains 60,000 images. How should you complete the body of the HTTP request? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content. NOTE: Each correct selection is worth one point.

Values	Answer Area
<input type="text" value="faceListId"/>	<pre>{ "faceId": "18c51a87-3a69-47a8-aedc-a54745f708a1", []: "employeefaces", "maxNumOfCandidatesReturned": 1, "mode": [] }</pre>
<input type="text" value="LargeFaceListId"/>	
<input type="text" value="matchFace"/>	
<input type="text" value="matchPerson"/>	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: LargeFaceListID
 LargeFaceList: Add a face to a specified large face list, up to 1,000,000 faces. Note: Given query face's faceId, to search the similar-looking faces from a faceId array, a face list or a large face list. A "faceListId" is created by FaceList - Create containing persistedFaceIds that will not expire. And a "largeFaceListId" is created by LargeFaceList - Create containing persistedFaceIds that will also not expire.

NEW QUESTION 50

- (Topic 3)
 You have the following C# method for creating Azure Cognitive Services resources programmatically.

```
static void create_resource(CognitiveServicesManagementClient client, string resource_name, string kind, string account_tier, string location)
{
    CognitiveServicesAccount parameters =
        new CognitiveServicesAccount(null, null, kind, location, resource_name,
        new CognitiveServicesAccountProperties(), new Sku(account_tier));
    var result = client.Accounts.Create(resource_group_name, account_tier,
parameters);
}
```

You need to call the method to create a free Azure resource in the West US Azure region. The resource will be used to generate captions of images automatically. Which code should you use?

- A. `create_resource(client, "res1", "ComputerVision", "F0", "westus")`
- B. `create_resource(client, "res1", "CustomVision.Prediction", "F0", "westus")`
- C. `create_resource(client, "res1", "ComputerVision", "S0", "westus")`
- D. `create_resource(client, "res1", "CustomVision.Prediction", "S0", "westus")`

Answer: B

Explanation:

<https://azure.microsoft.com/en-us/pricing/details/cognitive-services/computer-vision/>

NEW QUESTION 55

- (Topic 3)

You have a Language service resource that performs the following:

- Sentiment analysis
- Named Entity Recognition (NER)
- Personally Identifiable Information (PII) identification

You need to prevent the resource from persisting input data once the data is analyzed. Which query parameter in the Language service API should you configure?

- A. `loggingOptOut`
- B. `piiCategories`
- C. `showStats`
- D. `Model-version`

Answer: A

NEW QUESTION 59

- (Topic 3)

You are developing a new sales system that will process the video and text from a public-facing website.

You plan to monitor the sales system to ensure that it provides equitable results regardless of the user's location or background.

Which two responsible AI principles provide guidance to meet the monitoring requirements? Each correct answer presents part of the solution. (Choose two.)

NOTE: Each correct selection is worth one point.

- A. transparency
- B. fairness
- C. inclusiveness
- D. reliability and safety
- E. privacy and security

Answer: BC

Explanation:

<https://docs.microsoft.com/en-us/learn/modules/get-started-ai-fundamentals/8-understand-responsible-ai>

NEW QUESTION 60

- (Topic 3)

You are building an app that will share user images.

You need to configure the app to perform the following actions when a user uploads an image:

- Categorize the image as either a photograph or a drawing.
- Generate a caption for the image.

The solution must minimize development effort.

Which two services should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. object detection in Computer Vision
- B. content tags in Computer Vision
- C. image descriptions in Computer Vision
- D. image type detection in Computer Vision
- E. image classification in Custom Vision

Answer: CD

Explanation:

? According to the Microsoft documentation, Computer Vision is a cloud-based service that provides developers with access to advanced algorithms for processing images and returning information. By uploading an image or specifying an image URL, Computer Vision algorithms can analyze visual content in different ways based on inputs and user choices.

? According to the Microsoft documentation, image type detection is one of the features of Computer Vision that can categorize an image as either a photograph or a drawing. You can use the image type detection feature by calling the Analyze Image API with the `visualFeatures` parameter set to `ImageType`. The API will return a JSON response with an `imageType` field that indicates whether the image is a photo or a clipart.

? According to the Microsoft documentation, image descriptions is another feature of Computer Vision that can generate a caption for an image. You can use the image descriptions feature by calling the Analyze Image API with the `visualFeatures` parameter set to `Description`. The API will return a JSON response with a `description` field that contains a list of captions for the image, each with a confidence score.

? Therefore, by using these two features of Computer Vision, you can achieve your app requirements with minimal development effort. You don't need to use any other services, such as object detection, content tags, or Custom Vision, which are designed for different purposes.

NEW QUESTION 61

- (Topic 3)

You have an Azure subscription that contains an Azure App Service app named App1. You provision a multi-service Azure Cognitive Services resource named CSAccount1.

You need to configure App1 to access CSAccount1. The solution must minimize administrative effort.

What should you use to configure App1?

- A. a system-assigned managed identity and an X.509 certificate
- B. the endpoint URI and an OAuth token
- C. the endpoint URI and a shared access signature (SAS) token
- D. the endpoint URI and subscription key

Answer: D

NEW QUESTION 66

DRAG DROP - (Topic 3)

You are building a customer support chatbot.

You need to configure the bot to identify the following:

- Code names for internal product development
- Messages that include credit card numbers

The solution must minimize development effort. Which Azure Cognitive Service for Language feature should you use for each requirement? To answer, drag the appropriate features to the correct requirements. Each feature may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content

NOTE: Each correct selection is worth one point.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 71

HOTSPOT - (Topic 3)

You have a collection of press releases stored as PDF files.

You need to extract text from the files and perform sentiment analysis.

Which service should you use for each task? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.


Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 74

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area




- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 77

DRAG DROP - (Topic 3)

You train a Custom Vision model to identify a company's products by using the Retail domain.

You plan to deploy the model as part of an app for Android phones. You need to prepare the model for deployment.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

- Change the model domain.
- Retrain the model.
- Test the model.
- Export the model.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

In user want to change to deploy offline model

- * 1. Change model domain to compact model
- * 2. Retrain compact model
- * 3. Export model

NEW QUESTION 82

- (Topic 3)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
 You create a web app named app1 that runs on an Azure virtual machine named vm1. Vm1 is on an Azure virtual network named vnet1.
 You plan to create a new Azure Cognitive Search service named service1.
 You need to ensure that app1 can connect directly to service1 without routing traffic over the public internet.
 Solution: You deploy service1 and a public endpoint, and you configure a network security group (NSG) for vnet1.
 Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:


<https://docs.microsoft.com/en-us/azure/virtual-network/network-security-groups-overview#network-security-groups>
<https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-overview>

NEW QUESTION 86

HOTSPOT - (Topic 3)

You have a library that contains thousands of images.
 You need to tag the images as photographs, drawings, or clipart.
 Which service endpoint and response property should you use? To answer, select the appropriate options in the answer area.
 NOTE: Each correct selection is worth one point.

Answer Area




- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 91

HOTSPOT - (Topic 3)

You are building a chatbot by using the Microsoft Bot Framework SDK.
 You use an object named UserProfile to store user profile information and an object named ConversationData to store information related to a conversation.
 You create the following state accessors to store both objects in state.

```
var userStateAccessors = _userState.CreateProperty<UserProfile>(nameof(UserProfile)); var conversationStateAccessors =
_conversationState.CreateProperty<ConversationData>(nameof(ConversationData));
```

 The state storage mechanism is set to Memory Storage.
 For each of the following statements, select Yes if the statement is true. Otherwise, select No.
 NOTE: Each correct selection is worth one point.

Statements	Yes	No
The code will create and maintain the <code>UserProfile</code> object in the underlying storage layer.	<input type="checkbox"/>	<input type="checkbox"/>
The code will create and maintain the <code>ConversationData</code> object in the underlying storage layer.	<input type="checkbox"/>	<input type="checkbox"/>
The <code>UserProfile</code> and <code>ConversationData</code> objects will persist when the Bot Framework runtime terminates.	<input type="checkbox"/>	<input type="checkbox"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

You create property accessors using the `CreateProperty` method that provides a handle to the `BotState` object. Each state property accessor allows you to get or set the value of the associated state property.

Box 2: Yes

Box 3: No

Before you exit the turn handler, you use the state management objects' `SaveChangesAsync()` method to write all state changes back to storage.

NEW QUESTION 92

- (Topic 3)

You have a chatbot that was built by using Microsoft Bot Framework and deployed to Azure.

You need to configure the bot to support voice interactions. The solution must support multiple client apps.

Which type of channel should you use?

- A. Cortana
- B. Microsoft Teams
- C. Direct Line Speech

Answer: C

NEW QUESTION 95

- (Topic 3)

You build a chatbot that uses the Azure OpenAI GPT 3.5 model.

You need to improve the quality of the responses from the chatbot. The solution must minimize development effort.

What are two ways to achieve the goal? Each correct answer presents a complete solution. NOTE: Each correct answer is worth one point.

- A. Fine-tune the model.
- B. Provide grounding content.
- C. Add sample request/response pairs.
- D. Retrain the language model by using your own data.
- E. Train a custom large language model (LLM).

Answer: BC

NEW QUESTION 98

- (Topic 3)

You are developing a method that uses the Computer Vision client library. The method will perform optical character recognition (OCR) in images. The method has the following code.

```
public static async Task ReadFileUrl(ComputerVisionClient client, string urlFile)
{
    const int numberOfCharsInOperationId = 36;

    var txtHeaders = await client.ReadAsync(urlFile, language: "en");

    string opLocation = txtHeaders.OperationLocation;
    string operationId = opLocation.Substring(opLocation.Length -
        numberOfCharsInOperationId);

    ReadOperationResult results;

    results = await client.GetReadResultAsync(Guid.Parse(operationId));

    var textUrlFileResults = results.AnalyzeResult.ReadResults;
    foreach (ReadResult page in textUrlFileResults)
    {
        foreach (Line line in page.Lines)
        {
            Console.WriteLine(line.Text);
        }
    }
}
```

During testing, you discover that the call to the GetReadResultAsync method occurs before the read operation is complete. You need to prevent the GetReadResultAsync method from proceeding until the read operation is complete. Which two actions should you perform? Each correct answer presents part of the solution. (Choose two.)
 NOTE: Each correct selection is worth one point.

- A. Remove the Guid.Parse(operationId) parameter.
- B. Add code to verify the results.Status value.
- C. Add code to verify the status of the txtHeaders.Status value.
- D. Wrap the call to GetReadResultAsync within a loop that contains a delay.

Answer: BD

Explanation:

Example code : do

```
{
    results = await client.GetReadResultAsync(Guid.Parse(operationId));
}
```

while ((results.Status == OperationStatusCodes.Running || results.Status == OperationStatusCodes.NotStarted));

Reference:

<https://github.com/Azure-Samples/cognitive-services-quickstart-code/blob/master/dotnet/ComputerVision/ComputerVisionQuickstart.cs>

NEW QUESTION 102

- (Topic 3)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have an Azure Cognitive Search service.

During the past 12 months, query volume steadily increased.

You discover that some search query requests to the Cognitive Search service are being throttled.

You need to reduce the likelihood that search query requests are throttled. Solution: You migrate to a Cognitive Search service that uses a higher tier. Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

A simple fix to most throttling issues is to throw more resources at the search service (typically replicas for query-based throttling, or partitions for indexing-based throttling). However, increasing replicas or partitions adds cost, which is why it is important to know the reason why throttling is occurring at all.

Reference:

<https://docs.microsoft.com/en-us/azure/search/search-performance-analysis>

NEW QUESTION 103

- (Topic 3)

You are developing an application that will use Azure Cognitive Search for internal documents.

You need to implement document-level filtering for Azure Cognitive Search.

Which three actions should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Send Azure AD access tokens with the search request.
- B. Retrieve all the groups.
- C. Retrieve the group memberships of the user.
- D. Add allowed groups to each index entry.
- E. Create one index per group.
- F. Supply the groups as a filter for the search requests.

Answer: CDF

Explanation:

Your documents must include a field specifying which groups have access. This information becomes the filter criteria against which documents are selected or rejected from the result set returned to the issuer.
 D: A query request targets the documents collection of a single index on a search service. CF: In order to trim documents based on group_ids access, you should issue a search query with a group_ids/any(g:search.in(g, 'group_id1, group_id2,...')) filter, where 'group_id1, group_id2,...' are the groups to which the search request issuer belongs.
 Reference:
<https://docs.microsoft.com/en-us/azure/search/search-security-trimming-for-azure-search>

NEW QUESTION 107

HOTSPOT - (Topic 3)

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Azure Databricks is an Apache Spark-based analytics platform.	<input type="radio"/>	<input type="radio"/>
Azure Analysis Services is used for transactional workloads.	<input type="radio"/>	<input type="radio"/>
Azure Data Factory orchestrates data integration workflows.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
Azure Databricks is an Apache Spark-based analytics platform.	<input checked="" type="radio"/>	<input type="radio"/>
Azure Analysis Services is used for transactional workloads.	<input checked="" type="radio"/>	<input type="radio"/>
Azure Data Factory orchestrates data integration workflows.	<input checked="" type="checkbox"/>	<input type="radio"/>

NEW QUESTION 111

DRAG DROP - (Topic 3)

You are developing a photo application that will find photos of a person based on a sample image by using the Face API.

You need to create a POST request to find the photos.

How should you complete the request? To answer, drag the appropriate values to the correct targets. Each value may be used once, more than once, or not at all.

You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Values

- detect
- findsimilars
- group
- identify
- matchFace
- matchPerson
- verify

Answer Area

```
POST {Endpoint}/face/v1.0/ 
Request Body

{
  "faceId": "c5c24a82-6845-4031-9d5d-978df9175426",
  "largeFaceListId": "sample_list",
  "largeFaceListId": "sample_list",
  "maxNumOfCandidatesReturned": 10,
  "mode": "
}
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: findsimilars

<https://docs.microsoft.com/en-us/rest/api/faceapi/face/find-similar>

Box 2: matchPerson

Find similar has two working modes, "matchPerson" and "matchFace". "matchPerson" is the default mode that it tries to find faces of the same person as possible by using internal same-person thresholds. It is useful to find a known person's other photos. Note that an empty list will be returned if no faces pass the internal thresholds. "matchFace" mode ignores same-person thresholds and returns ranked similar faces anyway, even the similarity is low. It can be used in the cases like searching celebrity-looking faces.

NEW QUESTION 115

- (Topic 3)

You are building a multilingual chatbot.

You need to send a different answer for positive and negative messages.

Which two Text Analytics APIs should you use? Each correct answer presents part of the solution. (Choose two.)
 NOTE: Each correct selection is worth one point.

- A. Linked entities from a well-known knowledge base
- B. Sentiment Analysis
- C. Key Phrases
- D. Detect Language
- E. Named Entity Recognition

Answer: BD

Explanation:

B: The Text Analytics API's Sentiment Analysis feature provides two ways for detecting positive and negative sentiment. If you send a Sentiment Analysis request, the API will return sentiment labels (such as "negative", "neutral" and "positive") and confidence scores at the sentence and document-level.
 D: The Language Detection feature of the Azure Text Analytics REST API evaluates text input for each document and returns language identifiers with a score that indicates the strength of the analysis.
 This capability is useful for content stores that collect arbitrary text, where language is unknown. Reference:
<https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-tos/text-analytics-how-to-sentiment-analysis?tabs=version-3-1>
<https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-tos/text-analytics-how-to-language-detection>

NEW QUESTION 118

DRAG DROP - (Topic 3)

You plan to use a Language Understanding application named app1 that is deployed to a container. App1 was developed by using a Language Understanding authoring resource named lu1. App1 has the versions shown in the following table.

Version	Trained date	Published date
V1.2	None	None
V1.1	2020-10-01	None
V1.0	2020-09-01	2020-09-15

You need to create a container that uses the latest deployable version of app1. Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order. (Choose three.)

Actions

Answer Area

- Run a container that has version set as an environment variable.
- Export the model by using the Export as JSON option.
- Select v1.1 of app1.
- Run a container and mount the model file.
- Select v1.0 of app1.
- Export the model by using the Export for containers (GZIP) option.
- Select v1.2 of app1.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Export the model using the Export for containers (GZIP) option. Export versioned app's package from LUIS portal
 The versioned app's package is available from the Versions list page.
 ? Sign on to the LUIS portal.
 ? Select the app in the list.
 ? Select Manage in the app's navigation bar.
 ? Select Versions in the left navigation bar.
 ? Select the checkbox to the left of the version name in the list.
 ? Select the Export item from the contextual toolbar above the list.
 ? Select Export for container (GZIP).
 ? The package is downloaded from the browser.



Step 2: Select v1.1 of app1.

A trained or published app packaged as a mounted input to the container with its associated App ID.

Step 3: Run a contain and mount the model file.

Run the container, with the required input mount and billing settings.

NEW QUESTION 122

- (Topic 3)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have a chatbot that uses question answering in Azure Cognitive Service for Language Users report that the responses of the chatbot lack formality when answering spurious questions

You need to ensure that the chatbot provides formal responses to spurious questions. Solution: From Language Studio, you remove all the chit-chat question and answer pairs, and then retrain and republish the model Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Removing all the chit-chat question and answer pairs from the project will not ensure that the chatbot provides formal responses to spurious questions. It will only make the chatbot unable to handle any chit-chat scenarios, which may result in a poor user experience and a loss of engagement. Instead, you should choose a chit-chat personality that matches the tone and style of your chatbot, such as Professional or Caring. You can also edit the chit-chat questions and answers to suit your specific needs, or add new ones that are not in the predefined data set¹². This way, you can ensure that the chatbot responds appropriately to spurious questions, while still maintaining a conversational and engaging interaction with the user.

NEW QUESTION 126

HOTSPOT - (Topic 3)

You are building a model that will be used in an iOS app.

You have images of cats and dogs. Each image contains either a cat or a dog.

You need to use the Custom Vision service to detect whether the images is of a cat or a dog.

How should you configure the project in the Custom Vision portal? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Project Types: ▼

- Classification
- Object Detection

Classification Types: ▼

- Multiclass (Single tag per image)
- Multilabel (Multiple tags per image)

Domains: ▼

- Audit
- Food
- General
- General (compact)
- Landmarks
- Landmarks (compact)
- Retail
- Retail (compact)

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Classification

Box 2: Multiclass

A multiclass classification project is for classifying images into a set of tags, or target labels. An image can be assigned to one tag only.

Box 3: General

General: Optimized for a broad range of image classification tasks. If none of the other specific domains are appropriate, or if you're unsure of which domain to choose, select one of the General domains.

NEW QUESTION 127

- (Topic 3)

Your company has a reporting solution that has paginated reports. The reports query a dimensional model in a data warehouse. Which type of processing does the reporting solution use?

- A. Online Transaction Processing (OLTP)
- B. Online Analytical Processing (OLAP)
- C. batch processing
- D. stream processing

Answer: B

NEW QUESTION 128

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

_____ is responsible for identifying which business rules must be applied to the data of a company.

A data analyst
 A data engineer
 A data scientist

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

_____ is responsible for identifying which business rules must be applied to the data of a company.

A data analyst
 A data engineer
 A data scientist

NEW QUESTION 133

- (Topic 3)

You have an Azure subscription that contains an Azure Cognitive Service for Language resource. You need to identify the URL of the REST interface for the Language service. Which blade should you use in the Azure portal?

- A. Identity
- B. Keys and Endpoint
- C. Properties
- D. Networking

Answer: B

NEW QUESTION 136

- (Topic 3)

You need to create a visualization of running sales totals per quarter as shown in the following exhibit.



What should you create in Power BI Desktop;1

- A. a waterfall chart
- B. a ribbon chart
- C. a bar chart
- D. a decomposition tree

Answer: C

NEW QUESTION 137

DRAG DROP - (Topic 3)

You train a Custom Vision model used in a mobile app.

You receive 1,000 new images that do not have any associated data.

You need to use the images to retrain the model. The solution must minimize how long it takes to retrain the model.

Which three actions should you perform in the Custom Vision portal? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
Upload the images by category.	
Get suggested tags.	
Upload all the images.	⬅️
Group the images locally into category folders.	➡️
Review the suggestions and confirm the tags.	
Tag the images manually.	⬆️
	⬇️

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

1) upload all the images 2.) Get suggested tags 3.) Review the suggestions and confirm the tags

NEW QUESTION 141

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

The ALTER
JOIN
SET
WHERE clause can be used in Data Manipulation Language (DML) statements to specify the criteria that rows must match.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

The ALTER
JOIN
SET
WHERE clause can be used in Data Manipulation Language (DML) statements to specify the criteria that rows must match.

NEW QUESTION 142

HOTSPOT - (Topic 3)

You plan to deploy an Azure OpenAI resource by using an Azure Resource Manager (ARM) template.

You need to ensure that the resource can respond to 600 requests per minute.

How should you complete the template? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

```
{
  "type": "Microsoft.CognitiveServices/accounts/deployments",
  "apiVersion": "2023-05-01",
  "name": "arm-aoai-sample-resource/arm-je-std-deployment",
  "dependsOn": [
    "[resourceId('Microsoft.CognitiveServices/accounts', 'arm-aoai-sample-resource')]"
  ],
  "sku": {
    "name": "Standard",
    "capacity": 600,
    "count": 1,
    "maxValue": 100,
    "size": 600
  },
  "properties": {
    "model": {
      "format": "OpenAI",
      ...
    }
  }
}
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
 Answer Area

```
{
  "type": "Microsoft.CognitiveServices/accounts/deployments",
  "apiVersion": "2023-05-01",
  "name": "arm-aoai-sample-resource/arm-je-std-deployment",
  "dependsOn": [
    "[resourceId('Microsoft.CognitiveServices/accounts', 'arm-aoai-sample-resource')]"
  ],
  "sku": {
    "name": "Standard",
    "capacity": 600,
    "count": 1,
    "maxValue": 100,
    "size": 600
  },
  "properties": {
    "model": {
      "format": "OpenAI",
      ...
    }
  }
}
```

NEW QUESTION 147

- (Topic 3)

You build a Language Understanding model by using the Language Understanding portal. You export the model as a JSON file as shown in the following sample.

```
{
  "text": "average amount of rain by month at chicago last year",
  "intent": "Weather.CheckWeatherValue",
  "entities": [
    {
      "entity": "Weather.WeatherRange",
      "startPos": 0,
      "endPos": 6,
      "children": []
    },
    {
      "entity": "Weather.WeatherCondition",
      "startPos": 18,
      "endPos": 21,
      "children": []
    },
    {
      "entity": "Weather.Historic",
      "startPos": 23,
      "endPos": 30,
      "children": []
    }
  ]
}
```

To what does the Weather.Historic entity correspond in the utterance?

- A. by month
- B. chicago
- C. rain
- D. location

Answer: A

NEW QUESTION 149

- (Topic 3)

What is used to define a query in a stream processing jobs in Azure Stream Analytics?

- A. SQL
- B. XML
- C. YAML
- D. KOL

Answer: A

NEW QUESTION 150

HOTSPOT - (Topic 3)

You are developing a text processing solution. You develop the following method.

```
static void GetKeyPhrases(TextAnalyticsClient textAnalyticsClient, string text)
{
    var response = textAnalyticsClient.ExtractKeyPhrases(text);
    Console.WriteLine("Key phrases:");

    foreach (string keyphrase in response.Value)
    {
        Console.WriteLine($"{keyphrase}");
    }
}
```

You call the method by using the following code. GetKeyPhrases(textAnalyticsClient, "the cat sat on the mat");

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
The call will output key phrases from the input string to the console.	<input type="radio"/>	<input type="radio"/>
The output will contain the following words: the, cat, sat, on, and mat.	<input type="radio"/>	<input type="radio"/>
The output will contain the confidence level for key phrases.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

The Key Phrase Extraction API evaluates unstructured text, and for each JSON document, returns a list of key phrases.

Box 2: No

'the' is not a key phrase.

This capability is useful if you need to quickly identify the main points in a collection of documents. For example, given input text "The food was delicious and there were wonderful staff", the service returns the main talking points: "food" and "wonderful staff".

Box 3: No

Key phrase extraction does not have confidence levels.

NEW QUESTION 152

- (Topic 3)

You are developing a method for an application that uses the Translator API.

The method will receive the content of a webpage, and then translate the content into Greek (el). The result will also contain a transliteration that uses the Roman alphabet.

You need to create the URI for the call to the Translator API. You have the following URI. <https://api.cognitive.microsofttranslator.com/translate?api-version=3.0>

Which three additional query parameters should you include in the URI? Each correct answer presents part of the solution. (Choose three.)

NOTE: Each correct selection is worth one point.

- A. toScript=Cyrl
- B. from=el
- C. textType=html
- D. to=el
- E. textType=plain
- F. toScript=Latn

Answer: ADF

Explanation:

C: textType is an optional parameter. It defines whether the text being translated is plain text or HTML text (used for web pages).

D: to is a required parameter. It specifies the language of the output text. The target language must be one of the supported languages included in the translation scope.

F: toScript is an optional parameter. It specifies the script of the translated text. We use Latin (Roman alphabet) script.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/translator/reference/v3-0-translate>

NEW QUESTION 157

- (Topic 3)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have a chatbot that uses question answering in Azure Cognitive Service for Language. Users report that the responses of the chatbot lack formality when answering random

questions that are outside the scope of the knowledge base.

You need to ensure that the chatbot provides formal responses to these spurious questions.

Solution: From Language Studio, you modify the question and answer pairs for the custom intents, and then retrain and republish the model.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 160

DRAG DROP - (Topic 3)

You have a chatbot that uses a QnA Maker application.

You enable active learning for the knowledge base used by the QnA Maker application. You need to integrate user input into the model.

Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Answer Area

- Add a task to the Azure resource.
- Approve and reject suggestions.
- Publish the knowledge base.
- Modify the automation task logic app to run an Azure Resource Manager template that creates the Azure Cognitive Services resource.
- For the knowledge base, select Show active learning suggestions.
- Save and train the knowledge base.
- Select the properties of the Azure Cognitive Services resource.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: For the knowledge base, select Show active learning suggestions.

In order to see the suggested questions, on the Edit knowledge base page, select View Options, then select Show active learning suggestions.

Step 2: Approve and reject suggestions.

Each QnA pair suggests the new question alternatives with a check mark, , to accept the question or an x to reject the suggestions. Select the check mark to add the question.

Step 3: Save and train the knowledge base.

Select Save and Train to save the changes to the knowledge base. Step 4: Publish the knowledge base.

Select Publish to allow the changes to be available from the GenerateAnswer API.

When 5 or more similar queries are clustered, every 30 minutes, QnA Maker suggests the alternate questions for you to accept or reject.

NEW QUESTION 162

- (Topic 2)

You are developing the document processing workflow.

You need to identify which API endpoints to use to extract text from the financial documents. The solution must meet the document processing requirements.

Which two API endpoints should you identify? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. /vision/v3.2/read/analyzeResults
- B. /formrecognizer/v2.0/prebuilt/receipt/analyze
- C. /vision/v3.2/read/analyze
- D. /vision/v3.2/describe
- E. /formercognizer/v2.0/custom/models{modelId}/ analyze

Answer: BC

Explanation:

C: Analyze Receipt - Get Analyze Receipt Result.

Query the status and retrieve the result of an Analyze Receipt operation. Request URL: <https://{endpoint}/formrecognizer/v2.0-preview/prebuilt/receipt/analyzeResults/{resultId}>

E: POST {Endpoint}/vision/v3.2/read/analyze

Use this interface to get the result of a Read operation, employing the state-of-the-art Optical Character Recognition (OCR) algorithms optimized for text-heavy documents.

Scenario: Contoso plans to develop a document processing workflow to extract information automatically from PDFs and images of financial documents

? The document processing solution must be able to process standardized financial

documents that have the following characteristics:

- Contain fewer than 20 pages.
- Be formatted as PDF or JPEG files.
- Have a distinct standard for each office.

? *The document processing solution must be able to extract tables and text from the financial documents.

? The document processing solution must be able to extract information from receipt images.

Reference:

<https://westus2.dev.cognitive.microsoft.com/docs/services/form-recognizer-api-v2- preview/operations/GetAnalyzeReceiptResult>

<https://docs.microsoft.com/en-us/rest/api/computervision/3.1/read/read>

NEW QUESTION 166

HOTSPOT - (Topic 2)

You are developing the knowledgebase by using Azure Cognitive Search. You need to build a skill that will be used by indexers. How should you complete the code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

```

{
  "@odata.type": "#Microsoft.Skills.Text.EntityRecognitionSkill",
  "categories": [],
  "categories": [ "Email", "Persons", "Organizations" ],
  "categories": [ "Locations", "Persons", "Organizations" ],
  "minimumPrecision": 0.7,
  "inputs": [
    { "name": "text",
      "source": "/document/content" }
  ],
  "outputs": [
    { "name": "persons", "targetName": "people" },
    { "name": "locations", "targetName": "locations" },
    { "name": "organizations", "targetName": "organizations" },
    { "name": "entities" },
    { "name": "categories" },
    { "name": "namedEntities" }
  ]
}

```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: "categories": ["Locations", "Persons", "Organizations"], Locations, Persons, Organizations are in the outputs.

Scenario: Contoso plans to develop a searchable knowledgebase of all the intellectual property

Note: The categories parameter is an array of categories that should be extracted. Possible category types: "Person", "Location", "Organization", "Quantity", "Datetime", "URL", "Email". If no category is provided, all types are returned.

Box 2: {"name": " entities"}

The include wikis, so should include entities in the outputs.

Note: entities is an array of complex types that contains rich information about the entities extracted from text, with the following fields

name (the actual entity name. This represents a "normalized" form) wikipediaId

wikipediaLanguage

wikipediaUrl (a link to Wikipedia page for the entity) etc.

NEW QUESTION 168

- (Topic 2)

You are developing the knowledgebase by using Azure Cognitive Search. You need to process wiki content to meet the technical requirements. What should you include in the solution?

- A. an indexer for Azure Blob storage attached to a skillset that contains the language detection skill and the text translation skill
- B. an indexer for Azure Blob storage attached to a skillset that contains the language detection skill
- C. an indexer for Azure Cosmos DB attached to a skillset that contains the document extraction skill and the text translation skill
- D. an indexer for Azure Cosmos DB attached to a skillset that contains the language detection skill and the text translation skill

Answer: C

Explanation:

The wiki contains text in English, French and Portuguese.

Scenario: All planned projects must support English, French, and Portuguese.

The Document Extraction skill extracts content from a file within the enrichment pipeline. This allows you to take advantage of the document extraction step that normally happens before the skillset execution with files that may be generated by other skills.

Note: The Translator Text API will be used to determine the from language. The Language detection skill is not required.

Reference:

<https://docs.microsoft.com/en-us/azure/search/cognitive-search-skill-document-extraction> <https://docs.microsoft.com/en-us/azure/search/cognitive-search-skill-text-translation>

NEW QUESTION 169

- (Topic 2)

You need to develop an extract solution for the receipt images. The solution must meet the document processing requirements and the technical requirements.

You upload the receipt images to the From Recognizer API for analysis, and the API returns the following JSON.

```

"documentResults": [
  {
    "docType": "prebuilt:receipt",
    "pageRange": [
      1,
      1
    ],
    "fields": {
      "ReceiptType": {
        "type": "string",
        "valueString": "Itemized",
        "confidence": 0.672
      },
      "MerchantName": {
        "type": "string",
        "valueString": "Tailwind",
        "text": "Tailwind",
        "boundingBox": [],
        "page": 1,
        "confidence": 0.913,
        "elements": [
          "#/readResults/0/lines/0/words/0"
        ]
      }
    }
  },
  ...

```

Which expression should you use to trigger a manual review of the extracted information by a member of the Consultant-Bookkeeper group?

- A. documentResults.docType == "prebuilt:receipt"
- B. documentResults.fields.".confidence < 0.7
- C. documentResults.fields.ReceiptType.confidence > 0.7
- D. documentResults.fields.MerchantName.confidence < 0.7

Answer: D

Explanation:

Need to specify the field name, and then use < 0.7 to handle trigger if confidence score is less than 70%.
 Reference:
<https://docs.microsoft.com/en-us/azure/applied-ai-services/form-recognizer/api-v2-0/reference-sdk-api-v2-0>

NEW QUESTION 174

DRAG DROP - (Topic 2)

You are developing a solution for the Management-Bookkeepers group to meet the document processing requirements. The solution must contain the following components:

- ? A Form Recognizer resource
 - ? An Azure web app that hosts the Form Recognizer sample labeling tool
- The Management-Bookkeepers group needs to create a custom table extractor by using the sample labeling tool.

Which three actions should the Management-Bookkeepers group perform in sequence? To answer, move the appropriate cmdlets from the list of cmdlets to the answer area and arrange them in the correct order.

Actions	Answer Area
Train a custom model	
Label the sample documents	➤
Create a new project and load sample documents	➤
Create a composite model	

- A. Mastered
- B. Not Mastered

Answer: A


Explanation:

Step 1: Create a new project and load sample documents
 Create a new project. Projects store your configurations and settings.
 Step 2: Label the sample documents
 When you create or open a project, the main tag editor window opens.
 Step 3: Train a custom model. Finally, train a custom model.

NEW QUESTION 178

DRAG DROP - (Topic 1)

You are developing the smart e-commerce project.
 You need to design the skillset to include the contents of PDFs in searches.
 How should you complete the skillset design diagram? To answer, drag the appropriate services to the correct stages. Each service may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
 NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Azure Blob storage

At the start of the pipeline, you have unstructured text or non-text content (such as images, scanned documents, or JPEG files). Data must exist in an Azure data storage service that can be accessed by an indexer.

Box 2: Computer Vision API

Scenario: Provide users with the ability to search insight gained from the images, manuals, and videos associated with the products.

The Computer Vision Read API is Azure's latest OCR technology (learn what's new) that extracts printed text (in several languages), handwritten text (English only), digits, and currency symbols from images and multi-page PDF documents.

Box 3: Translator API

Scenario: Product descriptions, transcripts, and all text must be available in English, Spanish, and Portuguese.

Box 4: Azure Files

Scenario: Store all raw insight data that was generated, so the data can be processed later.

NEW QUESTION 179

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

AI-102 Practice Exam Features:

- * AI-102 Questions and Answers Updated Frequently
- * AI-102 Practice Questions Verified by Expert Senior Certified Staff
- * AI-102 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * AI-102 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AI-102 Practice Test Here](#)