



Microsoft

Exam Questions AI-102

Designing and Implementing an Azure AI Solution

NEW QUESTION 1

DRAG DROP - (Topic 3)

You are using a Language Understanding service to handle natural language input from the users of a web-based customer agent.

The users report that the agent frequently responds with the following generic response: "Sorry, I don't understand that."

You need to improve the ability of the agent to respond to requests.

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

- Add prebuilt domain models as required.
- Validate the utterances logged for review and modify the model.
- Migrate authoring to an Azure resource authoring key.
- Enable active learning.
- Enable log collection by using Log Analytics.
- Train and republish the Language Understanding model.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

- enable active learning
- validate the utterances
- train and republish

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-how-to-review-endpoint-utterances#log-user-queries-to-enable-active-learning>
<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-concept-prebuilt-model>


NEW QUESTION 2

DRAG DROP - (Topic 3)

Match the Azure services to the appropriate locations in the architecture.

To answer, drag the appropriate service from the column on the left to its location on the right. Each service may be used once, more than once, or not at all.


NOTE: Each correct match is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 3

- (Topic 3)

You have an Azure IoT hub that receives series data from machinery. You need to build an app that will perform the following actions:

- Perform anomaly detection across multiple correlated sensors
- Identify the root cause of process stops.
- Send incident alerts

The solution must minimize development time. Which Azure service should you use?

- A. Azure Metrics Advisor
- B. Form Recognizer
- C. Azure Machine learning
- D. Anomaly Detector

Answer: D

NEW QUESTION 4


HOTSPOT - (Topic 3)

You are building content for a video training solution.

You need to create narration to accompany the video content. The solution must use Custom Neural Voice.

What should you use to create a custom neural voice, and which service should you use to generate the narration? To answer, select the appropriate options in the answer area.


NOTE: Each correct answer is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 5

- (Topic 3)

You have an Azure subscription. The subscription contains an Azure OpenAI resource that hosts a GPT-4 model named Model1 and an app named App1. App1 uses Model1.

You need to ensure that App1 will NOT return answers that include hate speech. What should you configure for Model1?

- A. the Frequency penalty parameter
- B. abuse monitoring
- C. a content filter
- D. the Temperature parameter

Answer: B

NEW QUESTION 6

- (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 7

HOTSPOT - (Topic 3)

You are building a chatbot.

You need to use the Content Moderator service to identify messages that contain sexually explicit language.

Which section in the response from the service will contain the category score, and which category will be assigned to the message? 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 8

- (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 9

- (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 10

- (Topic 3)

You are building a Language Understanding solution.

You discover that many intents have similar utterances containing airport names or airport codes.

You need to minimize the number of utterances used to train the model. Which type of custom entity should you use?

- A. Pattera.any
- B. machine-learning
- C. list
- D. regular expression

Answer: C

NEW QUESTION 10

- (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 build a language model by using a Language Understanding service. The language model is used to search for information on a contact list by using an intent named FindContact.

A conversational expert provides you with the following list of phrases to use for training.

Find contacts in London. Who do I know in Seattle? Search for contacts in Ukraine.

You need to implement the phrase list in Language Understanding. Solution: You create a new pattern in the FindContact intent.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Instead use a new intent for location.

Note: An intent represents a task or action the user wants to perform. It is a purpose or goal expressed in a user's utterance.

Define a set of intents that corresponds to actions users want to take in your application. Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-concept-intent>

NEW QUESTION 13

- (Topic 3)

You build a conversational bot named bot1.

You need to configure the bot to use a QnA Maker application.

From the Azure Portal, where can you find the information required by bot1 to connect to the QnA Maker application?

- A. Access control (IAM)
- B. Properties
- C. Keys and Endpoint
- D. Identity

Answer: C

Explanation:

Obtain values to connect your bot to the knowledge base 1. In the QnA Maker site, select your knowledge base. 2. With your knowledge base open, select the SETTINGS tab. Record the value shown for service name. This value is useful for finding your knowledge base of interest when using the QnA Maker portal interface. It's not used to connect your bot app to this knowledge base. 3. Scroll down to find Deployment details and record the following values from the Postman sample HTTP request: 4. POST /knowledgebases/<knowledge-base-id>/generateAnswer 5. Host: <your-host-url> 6. Authorization: EndpointKey <your-endpoint-key>

Reference:

<https://docs.microsoft.com/en-us/azure/bot-service/bot-builder-howto-qna>

NEW QUESTION 18

HOTSPOT - (Topic 3)

You are building a bot and that will use Language Understanding. You have a LUDown file that contains the following content.

```
## Confirm
- confirm
- ok
- yes

## ExtractName
- call me steve !
- i am anna
- (i'm|i am) {@PersonName.Any}[.]
- my name is {@PersonName.Any}[.]

## Logout
- forget me
- log out

## SelectItem
- choose last
- choose the {@DirectionalReference=bottom left}
- choose {@DirectionalReference=top right}
- i like {@DirectionalReference=left} one

## SelectNone
- none

@ ml DirectionalReference
@ prebuilt personName
```

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

SelectItem is [answer choice].

a domain
an entity
an intent
an utterance

Choose {@DirectionalReference=top right} is [answer choice].

a domain
an entity
an intent
an utterance

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

SelectItem is [answer choice].

a domain
an entity
an intent
an utterance

Choose {@DirectionalReference=top right} is [answer choice].

a domain
an entity
an intent
an utterance

NEW QUESTION 23

- (Topic 3)

Which database transaction property ensures that individual transactions are executed only once and either succeed in their entirety or roll back?

- A. consistency
- B. isolation
- C. atomicity
- D. durability

Answer: A

NEW QUESTION 27

- (Topic 3)

You are building an app that will use the Azure Video Indexer service. You plan to train a language model to recognize industry-specific terms. You need to upload a file that contains the industry-specific terms. Which file format should you use?

- A. PDF
- B. XML
- C. TXT
- D. XLS

Answer: C

NEW QUESTION 29

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 33

- (Topic 3)

You are building an AI solution that will use Sentiment Analysis results from surveys to calculate bonuses for customer service staff. You need to ensure that the solution meets the Microsoft responsible AI principles. What should you do?

- A. Add a human review and approval step before making decisions that affect the staffs financial situation
- B. Include the Sentiment Analysis results when surveys return a low confidence score.
- C. Use all the surveys, including surveys by customers who requested that their account be deleted and their data be removed.
- D. Publish the raw survey data to a central location and provide the staff with access to the location.

Answer: A

NEW QUESTION 36

- (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 38

- (Topic 3)

You train a Conversational Language Understanding model to understand the natural language input of users. You need to evaluate the accuracy of the model before deploying it.

What are two methods you can use? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. From the language authoring REST endpoint, retrieve the model evaluation summary.
- B. From Language Studio, enable Active Learning, and then validate the utterances logged for review.
- C. From Language Studio, select Model performance.
- D. From the Azure portal, enable log collection in Log Analytics, and then analyze the logs.

Answer: AC

NEW QUESTION 40

- (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 43

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
    {
        IntentRecognizer,
        SpeakerRecognizer,
        SpeechSynthesizer,
        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 48

DRAG DROP - (Topic 3)

You have a Custom Vision resource named acvdev in a development environment. You have a Custom Vision resource named acvprod in a production environment.

In acvdev, you build an object detection model named obj1 in a project named proj1. You need to move obj1 to acvprod.

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	Answer Area
Use the ExportProject endpoint on acvdev.	
Use the GetProjects endpoint on acvdev.	
Use the ImportProject endpoint on acvprod.	⬅️
Use the ExportIteration endpoint on acvdev.	➡️
Use the GetIterations endpoint on acvdev.	
Use the UpdateProject endpoint on acvprod.	⬆️

A. Mastered
 B. Not Mastered

Answer: A

Explanation:

Actions	Answer Area
Use the ExportProject endpoint on acvdev.	
Use the GetProjects endpoint on acvdev.	
Use the ImportProject endpoint on acvprod.	⬅️
Use the ExportIteration endpoint on acvdev.	
Use the GetIterations endpoint on acvdev.	
Use the UpdateProject endpoint on acvprod.	⬆️

Answer Area
Use the GetProjects endpoint on acvdev.
Use the ExportProject endpoint on acvdev.
Use the ImportProject endpoint on acvprod.

NEW QUESTION 49

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

A JSON document is an example of

graph data.
 relational data.
 semi-structured data.
 unstructured data.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

A JSON document is an example of

graph data. — →
 relational data. →
 semi-structured data. ↗
 unstructured data.

NEW QUESTION 54

- (Topic 3)

You create a bot by using the Microsoft Bot Framework SDK.

You need to configure the bot to respond to events by using custom text responses. What should you use?

- A. an adaptive card
- B. an activity handler
- C. a dialog
- D. a skill

Answer: B

NEW QUESTION 59

- (Topic 3)

Which property of a transactional workload guarantees that each transaction is treated as a single unit that either succeeds completely or fails completely?

- A. isolation
- B. atomicity
- C. consistency
- D. durability

Answer: B

NEW QUESTION 63

- (Topic 3)

What is a characteristic of a non-relational database?

- A. full support for Transact-SQL
- B. a fixed schema
- C. self describing entities

Answer: C

NEW QUESTION 65

HOTSPOT - (Topic 3)

You are developing an application to recognize employees' faces by using the Face Recognition API. Images of the faces will be accessible from a URI endpoint. The application has the following code.

```
static async void AddFace(string subscription_key, string personGroupId, string personId, string imageURI)
{
    var client = new HttpClient();
    client.DefaultRequestHeaders.Add("Ocp-Apim-Subscription-Key", subscription_key);
    var endpointURI = $"https://westus.api.cognitive.microsoft.com/face/v1.0/persongroups/{personGroupId}/persons/{personId}/persistedFaces";
    HttpResponseMessage response;
    var body = "{ \"url\": \"" + imageURI + "\"}";
    var content = new StringContent(body, Encoding.UTF8, "application/json");
    var response = await client.PutAsync(endpointURI, content);
}
```

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 add a face image to a person object in a person group.	<input type="radio"/>	<input type="radio"/>
The code will work for a group of 10,000 people.	<input type="radio"/>	<input type="radio"/>
AddFace can be called multiple times to add multiple face images to a person object.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

B: see this example code from documentation that uses PersonGroup of size 10,000 : <https://docs.microsoft.com/en-us/azure/cognitive-services/face/face-api-how-to-topics/how-to-add-faces>
the question wants to trick you into thinking you need to use a LargePersonGroup for a size of 10,000 - but the documentation for it doesn't include this limitation or criteria: <https://docs.microsoft.com/en-us/azure/cognitive-services/face/face-api-how-to-topics/how-to-use-large-scale>

NEW QUESTION 68

- (Topic 3)

You have receipts that are accessible from a URL.

You need to extract data from the receipts by using Form Recognizer and the SDK. The solution must use a prebuilt model.

Which client and method should you use?

- A. the FormRecognizerClient client and the StartRecognizeReceiptsFromUri method
- B. the FormTrainingClient client and the StartRecognizeContentFromUri method
- C. the FormRecognizerClient client and the StartRecognizeReceiptsFromUri method
- D. the FormTrainingClient client and the StartRecognizeReceiptsFromUri method

Answer: D

Explanation:

To analyze receipts from a URL, use the StartRecognizeReceiptsFromUri method Example code:

```
private static async Task AnalyzeReceipt( FormRecognizerClient recognizerClient, string receiptUri)
{
```

```
    RecognizedFormCollection receipts = await recognizerClient.StartRecognizeReceiptsFromUri(new Uri(receiptUri)).WaitForCompletionAsync();
```

Reference:

<https://docs.microsoft.com/en-us/azure/applied-ai-services/form-recognizer/quickstarts/client-library>

NEW QUESTION 72

- (Topic 3)

You are building a Language Understanding model for an e-commerce chatbot. Users can speak or type their billing address when prompted by the chatbot.

You need to construct an entity to capture billing addresses. Which entity type should you use?

- A. machine learned
- B. Regex
- C. list
- D. Pattern.any

Answer: B

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-concept-entity-types>

ML Entity with Structure

An ML entity can be composed of smaller sub-entities, each of which can have its own properties. For example, Address could have the following structure:

Address: 4567 Main Street, NY, 98052, USA Building Number: 4567

Street Name: Main Street State: NY

Zip Code: 98052 Country: USA

NEW QUESTION 73

HOTSPOT - (Topic 3)

You run the following command.

```
docker run --rm -it -p 5000:5000 --memory 10g --cpus 2 \
mcr.microsoft.com/azure-cognitive-services/textanalytics/sentiment \
Eula=accept \
Billing={ENDPOINT_URI} \
ApiKey={API_KEY}
```

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

Statements	Yes	No
Going to <code>http://localhost:5000/status</code> will query the Azure endpoint to verify whether the API key used to start the container is valid.	<input type="radio"/>	<input type="radio"/>
The container logging provider will write log data.	<input type="radio"/>	<input type="radio"/>
Going to <code>http://localhost:5000/swagger</code> will provide the details to access the documentation for the available endpoints.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

`http://localhost:5000/status` : Also requested with GET, this verifies if the api-key used to start the container is valid without causing an endpoint query.

Box 2: Yes

The command saves container and LUIS logs to output mount at `C:\output`, located on container host

Box 3: Yes

`http://localhost:5000/swagger` : The container provides a full set of documentation for the endpoints and a Try it out feature. With this feature, you can enter your settings into a web-

based HTML form and make the query without having to write any code. After the query returns, an example CURL command is provided to demonstrate the HTTP headers and body format that's required.

NEW QUESTION 76

HOTSPOT - (Topic 3)

You develop a test method to verify the results retrieved from a call to the Computer Vision API. The call is used to analyze the existence of company logos in images. The call returns a collection of brands named `brands`.

You have the following code segment.

```
foreach (var brand in brands)
{
    if (brand.Confidence >= .75)
        Console.WriteLine($"Logo of {brand.Name} between {brand.Rectangle.X},
{brand.Rectangle.Y} and {brand.Rectangle.W}, {brand.Rectangle.H}");
}
```

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 return the name of each detected brand with a confidence equal to or higher than 75 percent.	<input type="radio"/>	<input type="radio"/>
The code will return coordinates for the bottom-left corner of the rectangle that contains the brand logo of the displayed brands.	<input type="radio"/>	<input type="radio"/>
The code will return coordinates for the bottom-right corner of the rectangle that contains the brand logo of the displayed brands.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Statements	Yes	No
The code will return the name of each detected brand with a confidence equal to or higher than 75 percent.	<input checked="" type="radio"/>	<input type="radio"/>
The code will return coordinates for the bottom-left corner of the rectangle that contains the brand logo of the displayed brands.	<input type="radio"/>	<input checked="" type="radio"/>
The code will return coordinates for the bottom-right corner of the rectangle that contains the brand logo of the displayed brands.	<input type="radio"/>	<input checked="" type="radio"/>

NEW QUESTION 77

DRAG DROP - (Topic 3)

You plan to build a chatbot to support task tracking.

You create a Conversational Language Understanding service named lu1.

You need to build a Conversational Language Understanding model to integrate into the chatbot. The solution must minimize development time to build 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.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 81

- (Topic 3)

You have a text-based chatbot.

You need to enable content moderation by using the Text Moderation API of Content Moderator. Which two service responses should you use? Each correct answer presents part of the solution NOTE: Each correct selection is worth one point.

- A. the adult classification score
- B. optical character recognition (OCR)
- C. personal data
- D. text classification
- E. the racy classification score

Answer: AD

NEW QUESTION 84

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 89

- (Topic 3)

You have an Azure Cognitive Search solution and a collection of blog posts that include a category field. You need to index the posts. The solution must meet the following requirements:

- Include the category field in the search results.
- Ensure that users can search for words in the category field.

• Ensure that users can perform drill down filtering based on category. Which index attributes should you configure for the category field?

- A. searchable, facetable, and retrievable
- B. retrievable, filterable, and sortable
- C. retrievable, facetable, and key
- D. searchable, sortable, and retrievable

Answer: B

NEW QUESTION 93

DRAG DROP - (Topic 3)

You have 100 chatbots that each has its own Language Understanding model. Frequently, you must add the same phrases to each model.

You need to programmatically update the Language Understanding models to include the new phrases.

How should you complete the code? 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
AddPhraseListAsync	var phraselistId = await client.Features. <input type="text"/>
Phraselist	(appId, versionId, new <input type="text"/>
PhraselistCreateObject	{
Phrases	EnabledForAllModels = false,
SavePhraselistAsync	IsExchangeable = true,
UploadPhraseListAsync	Name = "PL1",
	Phrases = "item1,item2,item3,item4,item5"
	});

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: AddPhraseListAsync Example: Add phraselist feature

```
var phraselistId = await client.Features.AddPhraseListAsync(appId, versionId, new
```

```
PhraselistCreateObject
```

```
{
EnabledForAllModels = false, IsExchangeable = true, Name = "QuantityPhraselist", Phrases = "few,more,extra"
});
```

Box 2: PhraselistCreateObject

NEW QUESTION 96

- (Topic 3)

You are building a chatbot.

You need to configure the chatbot to query a knowledge base. Which dialog class should you use?

- A. AdaptiveDialog
- B. QnAMakerDialog
- C. ComponentDialog
- D. SkillDialog

Answer: B

NEW QUESTION 97

- (Topic 3)

You are building a solution that will detect anomalies in sensor data from the previous 24 hours.

You need to ensure that the solution scans the entire dataset, at the same time, for anomalies.

Which type of detection should you use?

- A. batch
- B. streaming
- C. change point

Answer: A

Explanation:

Batch anomaly detection is a type of anomaly detection that scans the entire dataset at once for outliers and unusual patterns. Batch anomaly detection is suitable for offline analysis of historical data, such as sensor data from the previous 24 hours. Batch anomaly detection can use various techniques, such as statistical methods, machine learning methods, or hybrid methods, to identify anomalies in the data.

NEW QUESTION 98

- (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 add indexes. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Instead, you could migrate to a Cognitive Search service that uses a higher tier.

Note: 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 101

DRAG DROP - (Topic 3)

You have a monitoring solution that uses the Azure AI Anomaly Detector service. You provision a server named Server1 that has intermittent internet access. You need to deploy the Azure AI Anomaly Detector to Server 1.

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
From Server1, run the docker push command.	
Query the prediction endpoint of the Azure AI Anomaly Detector in Azure.	
Install the Docker Engine on Server1.	
From Server1, run the docker pull command.	
From Server1, run the docker run command.	
Query the prediction endpoint on Server1.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions	Answer Area
From Server1, run the docker push command.	
Query the prediction endpoint of the Azure AI Anomaly Detector in Azure.	
Install the Docker Engine on Server1.	Install the Docker Engine on Server1.
From Server1, run the docker pull command.	From Server1, run the docker pull command.
From Server1, run the docker run command.	From Server1, run the docker run command.
Query the prediction endpoint on Server1.	Query the prediction endpoint on Server1.

NEW QUESTION 104

DRAG DROP - (Topic 3)

You are building a Language Understanding model for purchasing tickets.

You have the following utterance for an intent named PurchaseAndSendTickets. Purchase [2 audit business] tickets to [Paris] [next Monday] and send tickets to [email@domain.com]

You need to select the entity types. The solution must use built-in entity types to minimize training data whenever possible.

Which entity type should you use for each label? To answer, drag the appropriate entity types to the correct labels. Each entity type 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.

Entity Types	Answer Area
Email	
List	Paris: <input type="text"/>
Regex	email@domain.com: <input type="text"/>
GeographyV2	2 audit business: <input type="text"/>
Machine learned	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: GeographyV2

The prebuilt geographyV2 entity detects places. Because this entity is already trained, you do not need to add example utterances containing GeographyV2 to the application intents.

Box 2: Email

Email prebuilt entity for a LUIS app: Email extraction includes the entire email address from an utterance. Because this entity is already trained, you do not need to add example utterances containing email to the application intents.

Box 3: Machine learned

The machine-learning entity is the preferred entity for building LUIS applications.

NEW QUESTION 106

- (Topic 3)

Which Azure Storage service implements the key/value model?

- A. Azure Files
- B. Azure Blob
- C. Azure Table
- D. Azure Queue

Answer: C

NEW QUESTION 110

DRAG DROP - (Topic 3)

You need to develop an automated call handling system that can respond to callers in their own language. The system will support only French and English. Which Azure Cognitive Services service should you use to meet each requirement? To answer, drag the appropriate services to the correct requirements. Each service may be used once, more than once, or not at all. You may need to drag the split bat between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Services	Answer Area
Speaker Recognition	
Speech to Text	
Text Analytics	Detect the incoming language: <input style="width: 150px; height: 20px;" type="text"/>
Text to Speech	Respond in the callers' own language: <input style="width: 150px; height: 20px;" type="text"/>
Translator	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Text Analytics

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.

Box 2: Translator

Translator is a cloud-based neural machine translation service that is part of the Azure Cognitive Services family of REST APIs. Translator can be used with any operating system and powers many Microsoft products and services used by thousands of businesses worldwide to perform language translation and other language-related operations.

NEW QUESTION 114

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


Extract text:	<div style="border: 1px solid gray; padding: 5px;"> <div style="background-color: #e0e0e0; padding: 2px;">Computer Vision</div> <div style="background-color: #e0e0e0; padding: 2px;">Azure Cognitive Search</div> <div style="background-color: #0070c0; color: white; padding: 2px;">Computer Vision</div> <div style="background-color: #e0e0e0; padding: 2px;">Form Recognizer</div> </div>
Perform sentiment analysis:	<div style="border: 1px solid gray; padding: 5px;"> <div style="background-color: #e0e0e0; padding: 2px;">Language</div> <div style="background-color: #e0e0e0; padding: 2px;">Azure Cognitive Search</div> <div style="background-color: #e0e0e0; padding: 2px;">Computer Vision</div> <div style="background-color: #e0e0e0; padding: 2px;">Form Recognizer</div> <div style="background-color: #0070c0; color: white; padding: 2px;">Language</div> </div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 118

- (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 build a language model by using a Language Understanding service. The language model is used to search for information on a contact list by using an intent named FindContact.

A conversational expert provides you with the following list of phrases to use for training. Find contacts in London.

Who do I know in Seattle? Search for contacts in Ukraine.

You need to implement the phrase list in Language Understanding. Solution: You create a new entity for the domain.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Instead use a new intent for location.

Note: An intent represents a task or action the user wants to perform. It is a purpose or goal expressed in a user's utterance.

Define a set of intents that corresponds to actions users want to take in your application. Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-concept-intent>

NEW QUESTION 119

- (Topic 3)

You are building a bot by using Microsoft Bot Framework.

You need to configure the bot to respond to spoken requests. The solution must minimize development effort.

What should you do?

- A. Deploy the bot to Azure and register the bot with a Direct Une Speech channel
- B. Integrate the bot with Cortana by using the Bot Framework SDK.
- C. Create an Azure function that will call the Speech service and connect the bot to the function.
- D. Deploy the bot to Azure and register the bot with a Microsoft Teams channel.

Answer: B

NEW QUESTION 124

HOTSPOT - (Topic 3)

You are building a chatbot by using the Microsoft Bot Framework Composer. You have the dialog design shown in the following exhibit.

AskForName > BeginDialog > Text
Show code

Prompt for text

Text input
Collection information - Ask for a word or sentence.
[Learn more](#)

Bot Asks
User Input
Other

Property ⓘ string ▾

user.name

Output Format ⓘ string ▾

ex. =toUpper(this.value), \$(toUpper(this.value))

Value ⓘ expression ▾

fx =coalesce(@user.Name.@personName)

Expected responses (intent: #TextInput_Response_GH5FTe)

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
user.name is an entity.	<input type="radio"/>	<input type="radio"/>
The dialog asks for a user name and a user age and assigns appropriate values to the user.name and user.age properties.	<input type="radio"/>	<input type="radio"/>
The chatbot attempts to take the first non-null entity value for userName or personName and assigns the value to user.name.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: No
 User.name is a property. Box 2: Yes
 Box 3: Yes
 The coalesce() function evaluates a list of expressions and returns the first non-null (or non-empty for string) expression.

NEW QUESTION 128

DRAG DROP - (Topic 3)

You have a Custom Vision service project that performs object detection. The project uses the General domain for classification and contains a trained model. You need to export the model for use on a network that is disconnected from the internet. 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 Domains to **General (compact)**.
- Retrain the model.
- Create a new classification model.
- Change the classification type.
- Export the model.

>
<

Answer Area

>
<

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 132

DRAG DROP - (Topic 3)

You have an Azure subscription that contains an Azure OpenAI resource named AH. You plan to build an app named App1 that will write press releases by using AM.

You need to deploy an Azure OpenAI model for App1. The solution must minimize development effort.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 136

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.

- 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 140

DRAG DROP - (Topic 3)

You Build a bot in JavaScript.

From the Azure Command-Line interface (CLI), you run the following command. az bot prepare-deploy

You need to deploy the bot to Azure.

Which three Azure CLI commands should you run in sequence? To answer, move the appropriate commands from the list of commands to the answer area and arrange them in the client order.

Commands

- az ad app credential
- az ad app create
- az deployment group create
- az webapp deployment source config-zip
- az ad app update

Answer Area

➤
➤

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

To deploy the bot to Azure, you should run the following three Azure CLI commands in sequence:

? az deployment group create This command will create the Azure resources for your bot using an ARM template and a parameters file. You need to specify the resource group name, the template file path, and the parameters file path. For example:

```
az deployment group create --resource-group myResourceGroup --template-file "deploymentTemplates\template-with-preexisting-rg.json" --parameters "deploymentTemplates\parameters-for-template-BotApp-with-rg.json"
```

This command will also output the app ID and password of your bot, which you will need for the next command1.

? az webapp deployment source config-zip This command will deploy your bot code

to the app service that you created in the previous step. You need to specify the resource group name, the app service name, and the zip file path of your bot code. For example:

```
az webapp deployment source config-zip --resource-group myResourceGroup --name myBotAppService --src "code.zip"
```

This command will also output the URL of your bot endpoint, which you will need for the next command2.

? az ad app update This command will update your bot registration with the endpoint

URL of your bot. You need to specify the app ID of your bot and the endpoint URL. For example:

```
az ad app update --id myBotAppId --set replyUrls="https://myBotAppService.azurewebsites.net/api/messages"
```

This command will complete the deployment process and make your bot ready to be tested3.

NEW QUESTION 143

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

When provisioning an Azure Cosmos DB _____, you need to specify which type of API you will use.

- account
- container
- database
- item

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

When provisioning an Azure Cosmos DB _____, you need to specify which type of API you will use.

- account
- container
- database
- item

NEW QUESTION 145

HOTSPOT - (Topic 3)

You are developing a streaming Speech to Text solution that will use the Speech SDK and MP3 encoding.

You need to develop a method to convert speech to text for streaming MP3 data. How should you complete the code? To answer, select the appropriate options in the

answer area.

NOTE: Each correct selection is worth one point.

```
var audioFormat = (AudioStreamContainerFormat.MP3);
var speechConfig = SpeechConfig.FromSubscription("18c51a87-3a69-47a8-aedc-a54745f708a1", "westus");
var audioConfig = AudioConfig.FromStreamInput(pushStream, audioFormat);
using (var recognizer = new (speechConfig, audioConfig))
{
    var result = await recognizer.RecognizeOnceAsync();
    var text = result.Text;
}
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application, email Description automatically generated

NEW QUESTION 146

HOTSPOT - (Topic 3)

You are building an app that will share user images.


You need to configure the app to meet the following requirement's

- uploaded images must be scanned and any text must be extracted from the images.
- Extracted text must be analyzed for the presence of profane language.
- The solution must minimize development effort.

what should you use for each requirement? To answer, select the apocopate opinions in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Extraction: b) Computer Vision

Computer Vision is a service that can analyze images and extract text from them using optical character recognition (OCR) or read API1. OCR can detect and extract printed or handwritten text from images in various languages and formats, such as PDF, TIFF, or JPEG2. Read API can perform asynchronous batch processing of text from multiple images, and can also handle text in tables, mixed languages, and rotated text3.

Profane Language Detection: c) Content Moderator

Content Moderator is a service that can detect and filter out potentially offensive or unwanted content in text, images, and videos. For text, Content Moderator can identify profane or abusive language, personal data, and custom terms that you define. You can also use Content Moderator to review and moderate the content manually or automatically using workflows.

NEW QUESTION 151

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 153

DRAG DROP - (Topic 3)

You have a Language Understanding solution that runs in a Docker container.

You download the Language Understanding container image from the Microsoft Container Registry (MCR).

You need to deploy the container image to a host computer.

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

- From the Language Understanding portal, retrain the model.
- From the host computer, run the container and specify the input directory.
- From the Language Understanding portal, export the solution as a package file.
- From the host computer, move the package file to the Docker input directory.
- From the host computer, build the container and specify the output directory.

Answer Area



- A. Mastered
- B. Not Mastered


Answer: A

Explanation:

Actions

- From the Language Understanding portal, retrain the model.
- From the host computer, run the container and specify the input directory.
- From the Language Understanding portal, export the solution as a package file.
- From the host computer, move the package file to the Docker input directory.
- From the host computer, build the container and specify the output directory.

Answer Area



NEW QUESTION 158

- (Topic 3)

You have an Azure Cognitive Search solution and an enrichment pipeline that performs Sentiment Analysis on social media posts.

You need to define a knowledge store that will include the social media posts and the Sentiment Analysis results.

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

NOTE: Each correct selection is worth one point.

- A. storageContainer
- B. tables
- C. storageConnectionString
- D. files
- E. objects

Answer: DE

NEW QUESTION 159

- (Topic 3)

You have an existing Azure Cognitive Search service.

You have an Azure Blob storage account that contains millions of scanned documents stored as images and PDFs.

You need to make the scanned documents available to search as quickly as possible. What should you do?

- A. Split the data into multiple blob container
- B. Create a Cognitive Search service for each container
- C. Within each indexer definition, schedule the same runtime execution pattern.
- D. Split the data into multiple virtual folder
- E. Create an indexer for each container
- F. Increase the search unit
- G. Within each indexer definition, schedule a sequential execution pattern.
- H. Create a Cognitive Search service for each type of document.
- I. Split the data into multiple virtual folder
- J. Create an indexer for each folder
- K. Increase the search unit
- L. Within each indexer definition, schedule the same runtime execution pattern.

Answer: D

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/search/search-howto-indexing-azure-blob-storage>

NEW QUESTION 164

- (Topic 3)

You are examining the Text Analytics output of an application.

The text analyzed is: "Our tour guide took us up the Space Needle during our trip to Seattle last week."

The response contains the data shown in the following table.

Text	Category	ConfidenceScore
Tour guide	PersonType	0.45
Space Needle	Location	0.38
Trip	Event	0.78
Seattle	Location	0.78
Last week	DateTime	0.80

Which Text Analytics API is used to analyze the text?

- A. Sentiment Analysis
- B. Named Entity Recognition
- C. Entity Linking
- D. Key Phrase Extraction

Answer: B

NEW QUESTION 167

- (Topic 3)

You need to implement a table projection to generate a physical expression of an Azure Cognitive Search index.

Which three properties should you specify in the skillset definition JSON configuration table node? Each correct answer presents part of the solution. (Choose three.)

NOTE: Each correct selection is worth one point.

- A. tableName
- B. generatedKeyName
- C. dataSource
- D. dataSourceConnection
- E. source

Answer: ABE

Explanation:

Defining a table projection.

Each table requires three properties:

tableName: The name of the table in Azure Storage.

generatedKeyName: The column name for the key that uniquely identifies this row. source: The node from the enrichment tree you are sourcing your enrichments from. This node is usually the output of a shaper, but could be the output of any of the skills.

Reference:

<https://docs.microsoft.com/en-us/azure/search/knowledge-store-projection-overview>

NEW QUESTION 170

HOTSPOT - (Topic 3)

A1 You have an Azure OpenAI resource named AH that hosts three deployments of the GPT 3.5 model. Each deployment is optimized for a unique workload. You plan to deploy three apps. Each app will access AM by using the REST API and will use the deployment that was optimized for the apps intended workload. You need to provide each app with access to AH and the appropriate deployment. The solution must ensure that only the apps can access AM. What should you use to provide access to AM. and what should each app use to connect to its appropriate deployment? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Provide access to AI1 by using:

- An API key
- A bearer token
- A shared access signature (SAS) token

Connect to the deployment by using:

- An API key
- A deployment endpoint
- A deployment name
- A deployment type

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Provide access to AI1 by using:

- An API key
- A bearer token
- A shared access signature (SAS) token

Connect to the deployment by using:

- An API key
- A deployment endpoint
- A deployment name
- A deployment type

NEW QUESTION 175

- (Topic 3)

You have a mobile app that manages printed forms.

You need the app to send images of the forms directly to Forms Recognizer to extract relevant information. For compliance reasons, the image files must not be stored in the cloud.

In which format should you send the images to the Form Recognizer API endpoint?

- A. raw image binary
- B. form URL encoded
- C. JSON

Answer: A

NEW QUESTION 177

- (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 182

- (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 = textHeaders.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 183

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 185

- (Topic 3)

What are two benefits of platform as a service (PaaS) relational database offerings in Azure, such as Azure SQL Database? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. reduced administrative effort for managing the server infrastructure
- B. complete control over backup and restore processes
- C. in-database machine learning services S3
- D. access to the latest features

Answer: AD

NEW QUESTION 189

- (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 190

- (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 192

- (Topic 3)

Which database transaction property ensures that transactional changes to a database are preserved during unexpected operating system restarts?

- A. durability
- B. atomicity
- C. consistency
- D. isolation

Answer: A

NEW QUESTION 197

- (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 Cover BI Desktop;1

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

Answer: C

NEW QUESTION 198

HOTSPOT - (Topic 3)

You have an Azure subscription that contains an Azure OpenAI resource named AH. You build a chatbot that will use AI1 to provide generative answers to specific questions. You need to ensure that the responses are more creative and less deterministic.

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

```

new ChatCompletionsOptions()
(
    Messages =
    {
        new ChatMessage(
            ChatRole.User, @"",
        ),
        Temperature = (float)1.0,
        PresencePenalty,
        TokenSelectionBiases
    },
    MaxTokens = 800,
);
    
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```

new ChatCompletionsOptions()
{
    Messages =
    {
        new ChatMessage(
            ChatRole.User, @""),
    },
    Temperature = (float)1.0,
    MaxTokens = 800,
});
    
```

NEW QUESTION 202

- (Topic 3)

You are building an Azure WebJob that will create knowledge bases from an array of URLs.

You instantiate a QnAMakerClient object that has the relevant API keys and assign the object to a variable named client.

You need to develop a method to create the knowledge bases.

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

NOTE: Each correct selection is worth one point.

- A. Create a list of FileDTO objects that represents data from the WebJob.
- B. Call the client
- C. Knowledgebas
- D. CreateAsync method.
- E. Create a list of QnADTO objects that represents data from the WebJob.
- F. Create a CreateKbDTO object.

Answer: AC

Explanation:

Reference:

<https://docs.microsoft.com/en-us/rest/api/cognitiveservices-qnamaker/qnamaker4.0/knowledgebase/create>

NEW QUESTION 206

HOTSPOT - (Topic 3)

You are building an app that will process incoming email and direct email messages to either French or English language support teams.

Which Azure Cognitive Services API should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

https://

api.cognitive.microsofttranslator.com	/text/analytics/v3.1/entities/recognition/general
api.cognitive.microsofttranslator.com	/text/analytics/v3.1/entities/recognition/general
eastus.api.cognitive.microsoft.com	/text/analytics/v3.1/languages
portal.azure.com	/translator/text/v3.0/translate?to=en
	/translator/text/v3.0/translate?to=fr

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

https://

api.cognitive.microsofttranslator.com	/text/analytics/v3.1/entities/recognition/general
api.cognitive.microsofttranslator.com	/text/analytics/v3.1/entities/recognition/general
eastus.api.cognitive.microsoft.com	/text/analytics/v3.1/languages
portal.azure.com	/translator/text/v3.0/translate?to=en
	/translator/text/v3.0/translate?to=fr

NEW QUESTION 210

- (Topic 3)

You have an Azure Cognitive Services model named Model that identifies the intent of text input. You develop an app in C# named App1. You need to configure App1 to use Model1. Which package should you add to App1?

- A. Azure.AI.Language.Conversations
- B. SpeechServicesToolkit
- C. Universal.Microsoft.CognitiveServices.Speech
- D. Xamarin.Cognitive.Speech

Answer: C

NEW QUESTION 212

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 214

- (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 217

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 218

- (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 223

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 227

- (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 230

- (Topic 2)

You are developing the knowledgebase.

You use Azure Video Analyzer for Media (previously Video indexer) to obtain transcripts of webinars.

You need to ensure that the solution meets the knowledgebase requirements. What should you do?

- A. Create a custom language model
- B. Configure audio indexing for videos only
- C. Enable multi-language detection for videos
- D. Build a custom Person model for webinar presenters

Answer: B

Explanation:

Can search content in different formats, including video

Audio and video insights (multi-channels). When indexing by one channel, partial result for those models will be available.

Keywords extraction: Extracts keywords from speech and visual text.

Named entities extraction: Extracts brands, locations, and people from speech and visual text via natural language processing (NLP).

NEW QUESTION 234

- (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 Form 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 236

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 238

HOTSPOT - (Topic 1)

You are planning the product creation project.
 You need to build the REST endpoint to create the multilingual product descriptions. How should you complete the URI? To answer, select the appropriate options in the answer area.
 NOTE: Each correct selection is worth one point.

Answer Area

api.cognitive.microsofttranslator.com	/detect	?api-version=3.0&to=es&to=pt
api-nam.cognitive.microsofttranslator.com	/languages	
westus.tts.speech.microsoft.com	/text-to-speech	
wwics.cognitiveservices.azure.com/translator	/translate	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: api-nam.cognitive.microsofttranslator.com
<https://docs.microsoft.com/en-us/azure/cognitive-services/translator/reference/v3-0-reference>
 Box 2: /translate

NEW QUESTION 241

HOTSPOT - (Topic 1)

You are developing the shopping on-the-go project.

You are configuring access to the QnA Maker resources.

Which role should you assign to AllUsers and LeadershipTeam? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

AllUsers:

▼
Cognitive Service User
Contributor
Owner
QnA Maker Editor
QnA Maker Read

LeadershipTeam:

▼
Cognitive Service User
Contributor
Owner
QnA Maker Editor
QnA Maker Read

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: QnA Maker Editor
 Scenario: Provide all employees with the ability to edit Q&As.
 The QnA Maker Editor (read/write) has the following permissions: Create KB API Update KB API Replace KB API Replace Alterations "Train API" [in new service model v5]
 Box 2: Contributor
 Scenario: Only senior managers must be able to publish updates. Contributor permission: All except ability to add new members to roles

NEW QUESTION 242

.....

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](#)