



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

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

Relational data is stored in

- a file system as unstructured data.
- a hierarchal folder structure.
- a tabular form of rows and columns.
- comma-separated value (CSV) files.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Relational data is stored in

- a file system as unstructured data.
- a hierarchal folder structure.
- a Tabular form of rows and columns.
- comma-separated value (CSV) files.

NEW QUESTION 3

- (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 4

- (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 5

HOTSPOT - (Topic 3)

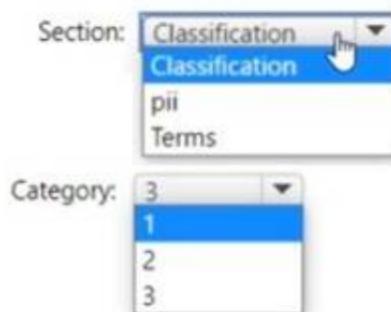
You are building a chatbot.

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

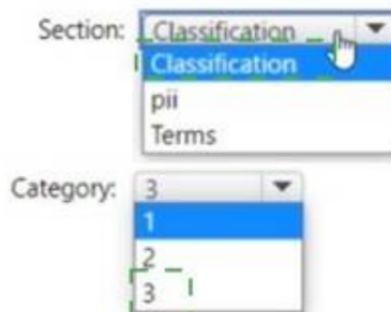


- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



NEW QUESTION 6

- (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 7

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

- A data analyst
- A data engineer
- A data scientist
- A database administrator

_____ is responsible for creating visuals and charts that help a company make informed decisions.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

- A data analyst
- A data engineer
- A data scientist
- A database administrator

_____ is responsible for creating visuals and charts that help a company make informed decisions.

NEW QUESTION 8

- (Topic 3)

You develop a custom question answering project in Azure Cognitive Service for Language. The project will be used by a chatbot. You need to configure the project to engage in multi-turn conversations. What should you do?

- A. Add follow-up prompts.
- B. Enable active learning.
- C. Add alternate questions.
- D. Enable chit-chat.

Answer: A

NEW QUESTION 9

- (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 10

DRAG DROP - (Topic 3)

You have an app that manages feedback.

You need to ensure that the app can detect negative comments by using the Sentiment Analysis API in Azure Cognitive Service for Language. The solution must ensure that the managed feedback remains on your company's internal network.

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.

NOTE: More than one order of answer choices is correct You will receive credit for any of the correct orders you select.

Actions

- Provision the Language service resource in Azure.
- Deploy a Docker container to an Azure container instance.
- Deploy a Docker container to an on-premises server.
- Identify the Language service endpoint URL and query the prediction endpoint.
- Run the container and query the prediction endpoint.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

- ? Provision the Language service resource in Azure.
- ? Deploy a Docker container to an on-premises server.
- ? Run the container and query the prediction endpoint.
- ? According to the Microsoft documentation, the Language service is a cloud-based service that provides various natural language processing features, such as

sentiment analysis, key phrase extraction, named entity recognition, etc. You can provision the Language service resource in Azure by following the steps in Create a Language resource. You will need to provide a name, a subscription, a resource group, a region, and a pricing tier for your resource. You will also get a key and an endpoint for your resource, which you will use to authenticate your requests to the Language service API.

? According to the Microsoft documentation, you can also use the Language service as a container on your own premises or in another cloud. This option gives you more control over your data and network, and allows you to use the Language service without an internet connection. You can deploy a Docker container to an on-premises server by following the steps in Deploy Language containers. You will need to have Docker installed on your server, pull the container image from the Microsoft Container Registry, and run the container with the appropriate parameters. You will also need to activate your container with your key and endpoint from your Azure resource.

? According to the Microsoft documentation, once you have deployed and activated your container, you can run it and query the prediction endpoint to get sentiment analysis results. The prediction endpoint is a local URL that follows this format: `http://<container IP address>:<port>/text/analytics/v3.1-preview.4/sentiment`. You can send HTTP POST requests to this endpoint with your text input in JSON format, and receive JSON responses with sentiment labels and scores for each document and sentence in your input.

NEW QUESTION 10

- (Topic 3)

You are examining the Language service 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 Language service API is used to analyze the Text?

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

Answer: B

NEW QUESTION 12

- (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 14

- (Topic 3)

You plan to build an app that will generate a list of tags for uploaded images. The app must meet the following requirements:

- Generate tags in a user's preferred language.
- Support English, French, and Spanish.
- Minimize development effort

You need to build a function that will generate the tags for the app. Which Azure service endpoint should you use?

- A. Custom Vision image classification
- B. Content Moderator Image Moderation
- C. Custom Translator
- D. Computer Vision Image Analysis

Answer: A

NEW QUESTION 15

- (Topic 3)

You are building a natural language model. You need to enable active learning.

What should you do?

- A. Add `show-all-intents=true` to the prediction endpoint query.
- B. Enable speech priming.
- C. Add `log=true` to the prediction endpoint query.
- D. Enable sentiment analysis.

Answer: C

Explanation:

Reference:

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

NEW QUESTION 16

- (Topic 3)

You deploy a web app that is used as a management portal for indexing in Azure Cognitive Search. The app is configured to use the primary admin key. During a security review, you discover unauthorized changes to the search index. You suspect that the primary access key is compromised. You need to prevent unauthorized access to the index management endpoint. The solution must minimize downtime. What should you do next?

- A. Regenerate the primary admin key, change the app to use the secondary admin key, and then regenerate the secondary admin key.
- B. Change the app to use a query key, and then regenerate the primary admin key and the secondary admin key.
- C. Regenerate the secondary admin key, change the app to use the secondary admin key, and then regenerate the primary key.
- D. Add a new query key, change the app to use the new query key, and then delete all the unused query keys.

Answer: C

NEW QUESTION 17

- (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 20

- (Topic 3)

You need to store event log data that is semi-structured and received as the logs occur. What should you use?

- A. Azure Table storage
- B. Azure Queue storage
- C. Azure Files

Answer: A

NEW QUESTION 23

- (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 24

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-afb4aa07d85", "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 25

- (Topic 3)

You have the following data sources:

- ? Finance: On-premises Microsoft SQL Server database
- ? Sales: Azure Cosmos DB using the Core (SQL) API
- ? Logs: Azure Table storage
- ? HR: Azure SQL database

You need to ensure that you can search all the data by using the Azure Cognitive Search REST API. What should you do?

- A. Configure multiple read replicas for the data in Sales.
- B. Mirror Finance to an Azure SQL database.
- C. Migrate the data in Sales to the MongoDB API.
- D. Ingest the data in Logs into Azure Sentinel.

Answer: B

Explanation:

On-premises Microsoft SQL Server database cannot be used as an index data source. Note: Indexer in Azure Cognitive Search: : Automate aspects of an indexing operation by configuring a data source and an indexer that you can schedule or run on demand. This feature is supported for a limited number of data source types on Azure. Indexers crawl data stores on Azure.

- ? Azure Blob Storage
- ? Azure Data Lake Storage Gen2 (in preview)
- ? Azure Table Storage
- ? Azure Cosmos DB
- ? Azure SQL Database
- ? SQL Managed Instance
- ? SQL Server on Azure Virtual Machines

Reference:

<https://docs.microsoft.com/en-us/azure/search/search-indexer-overview#supported-data-sources>

NEW QUESTION 28

- (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 develop an application to identify species of flowers by training a Custom Vision model. You receive images of new flower species.

You need to add the new images to the classifier.

Solution: You add the new images and labels to the existing model. You retrain the model, and then publish the model.

Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

The model needs to be extended and retrained.

NEW QUESTION 30

- (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 35

- (Topic 3)

You use the Custom Vision service to build a classifier.

After training is complete, you need to evaluate the classifier.

Which two metrics are available for review? Each correct answer presents a complete solution. (Choose two.)

NOTE: Each correct selection is worth one point.

- A. recall
- B. F-score
- C. weighted accuracy
- D. precision
- E. area under the curve (AUC)

Answer: AD

Explanation:

Custom Vision provides three metrics regarding the performance of your model: precision, recall, and AP.

Reference:

<https://www.tallan.com/blog/2020/05/19/azure-custom-vision/>

NEW QUESTION 38

- (Topic 3)

You plan to perform predictive maintenance.

You collect IoT sensor data from 100 industrial machines for a year. Each machine has 50 different sensors that generate data at one-minute intervals. In total, you have 5,000 time series datasets.

You need to identify unusual values in each time series to help predict machinery failures. Which Azure Cognitive Services service should you use?

- A. Anomaly Detector
- B. Cognitive Search
- C. Form Recognizer
- D. Custom Vision

Answer: A

NEW QUESTION 39

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

NEW QUESTION 44

HOTSPOT - (Topic 3)

You are building an Azure web app named App1 that will translate text from English to Spanish. You need to use the Text Translation REST API to perform the translation. The solution must ensure that you have data sovereignty in the United States. 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

The screenshot shows a URI completion interface. The URI is `https://` followed by a dropdown menu with the following options: `api-nam.cognitive.microsofttranslator.com`, `api.cognitive.microsofttranslator.com`, `api-nam.cognitive.microsofttranslator.com`, `api-nam.cognitiveservices.azure.com`, and `eastus.api.cognitive.microsoft.com`. The selected option is `api-nam.cognitive.microsofttranslator.com`. The URI continues with `/` followed by another dropdown menu with the following options: `translate`, `detect`, `languages`, `text-to-speech`, and `translate`. The selected option is `translate`. The URI ends with `?api-version=3.0&to=es`.

- A. Mastered

B. Not Mastered

Answer: A

Explanation:

* 1. api-nam.cognitive.microsofttranslator.com

* 2. translate

<https://learn.microsoft.com/en-us/azure/cognitive-services/Translator/reference/v3-0-reference#base-urls>

Requests to Translator are, in most cases, handled by the datacenter that is closest to where the request originated. If there's a datacenter failure when using the global endpoint, the request may be routed outside of the geography.

To force the request to be handled within a specific geography, use the desired geographical endpoint. All requests are processed among the datacenters within the geography.

- United States

api-nam.cognitive.microsofttranslator.com

<https://learn.microsoft.com/en-us/azure/cognitive-services/translator/reference/rest-api-guide-translate>

Translate specified source language text into the target language text.

NEW QUESTION 49

- (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 53

- (Topic 3)

You plan create an index for an Azure Cognitive Search service by using the Azure portal. The Cognitive Search service will connect to an Azure SQL database

The Azure SQL database contains a table named UserMessages. Each row m User

Messages has a field named MessageCopy that contains the text of social media messages sent by a user

Users win perform full text searches against the MessageCopy field, and the values of the field will be shown to the users-

You need to configure the properties of the index for the MessageCopy field to support the solution.

Winch attributes should you enable for the field?

- A. Searchable arc Retrievable
- B. Sortable and Retrievable
- C. Searchable arc Facetable
- D. Filterable and Retrievable

Answer: A

NEW QUESTION 56

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 http://localhost:5000/status 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 http://localhost:5000/swagger 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 57

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 60

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 63

HOTSPOT - (Topic 3)

You are building an app by using the Speech SDK. The app will translate speech from French to German by using natural language processing.

You need to define the source language and the output language.

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

```
var speechTranslationConfig =
SpeechTranslationConfig.FromSubscription(speechKey, speechRegion);
speechTranslationConfig.
speech_translation_config.
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```
var speechTranslationConfig =
SpeechTranslationConfig.FromSubscription(speechKey, speechRegion);
speechTranslationConfig.
speech_translation_config.
```

NEW QUESTION 65

- (Topic 3)

You are training a Language Understanding model for a user support system.

You create the first intent named GetContactDetails and add 200 examples. You need to decrease the likelihood of a false positive.

What should you do?

- A. Enable active learning.
- B. Add a machine learned entity.
- C. Add additional examples to the GetContactDetails intent.
- D. Add examples to the None intent.

Answer: A

Explanation:

Active learning is a technique of machine learning in which the machine learned model is used to identify informative new examples to label. In LUIS, active learning refers to adding utterances from the endpoint traffic whose current predictions are unclear to improve your model.

Reference:

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

NEW QUESTION 67

- (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, facettable, and retrievable
- B. retrievable, filterable, and sortable
- C. retrievable, facettable, and key
- D. searchable, sortable, and retrievable

Answer: B

NEW QUESTION 68

DRAG DROP - (Topic 3)

You are building an app that will scan confidential documents and use the Language service to analyze the contents.

You provision an Azure Cognitive Services resource.

You need to ensure that the app can make requests to the Language service endpoint. The solution must ensure that confidential documents remain on-premises. 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

- Pull an image from Docker Hub.
- Provision an on-premises Kubernetes cluster that has internet connectivity.
- Provision an Azure Kubernetes Service (AKS) resource.
- Run the container and specify an App ID and Client Secret.
- Provision an on-premises Kubernetes cluster that is isolated from the internet.
- Pull an image from the Microsoft Container Registry (MCR).
- Run the container and specify an API key and the Endpoint URL of the Cognitive Services resource.

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions

- Pull an image from Docker Hub.
- Provision an on-premises Kubernetes cluster that has internet connectivity.
- Provision an Azure Kubernetes Service (AKS) resource.
- Run the container and specify an App ID and Client Secret.
- Provision an on-premises Kubernetes cluster that is isolated from the internet.
- Pull an image from the Microsoft Container Registry (MCR).
- Run the container and specify an API key and the Endpoint URL of the Cognitive Services resource.

Answer Area

Provision an on-premises Kubernetes cluster that is isolated from the internet.

Pull an image from the Microsoft Container Registry (MCR).

Run the container and specify an API key and the Endpoint URL of the Cognitive Services resource.

NEW QUESTION 70

HOTSPOT - (Topic 3)

You are building an app that will enable users to upload images. The solution must meet the following requirements:

- Automatically suggest alt text for the images.
- Detect inappropriate images and block them.
- Minimize development effort.

You need to recommend a computer vision endpoint for each requirement.

What should you recommend? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Generate alt text:

- <https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image>
- <https://westus.api.cognitive.microsoft.com/contentmoderator/moderate/v1.0/ProcessImage/Evaluate>
- <https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image>
- <https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description>

Detect inappropriate content:

- <https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description>
- <https://westus.api.cognitive.microsoft.com/contentmoderator/moderate/v1.0/ProcessImage/Evaluate>
- <https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image>
- <https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description>
- <https://westus.api.cognitive.microsoft.com/vision/v3.2/describe?maxCandidates=1>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Generate alt text:

- <https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image>
- <https://westus.api.cognitive.microsoft.com/contentmoderator/moderate/v1.0/ProcessImage/Evaluate>
- <https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image>
- <https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description>

Detect inappropriate content:

- <https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description>
- <https://westus.api.cognitive.microsoft.com/contentmoderator/moderate/v1.0/ProcessImage/Evaluate>
- <https://westus.api.cognitive.microsoft.com/customvision/v3.1/prediction/projectId/classify/iterations/publishedName/image>
- <https://westus.api.cognitive.microsoft.com/vision/v3.2/analyze/?visualFeatures=Adult,Description>
- <https://westus.api.cognitive.microsoft.com/vision/v3.2/describe?maxCandidates=1>

NEW QUESTION 75

- (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 80

- (Topic 3)

You have an app that analyzes images by using the Computer Vision API.

You need to configure the app to provide an output for users who are vision impaired. The solution must provide the output in complete sentences.

Which API call should you perform?

- A. readInStreamAsync
- B. describeImageInStreamAsync
- C. toggleImageInStreamAsync
- D. analyzeImageByDomainInStreamAsync

Answer: D

NEW QUESTION 84

DRAG DROP - (Topic 3)

You are developing a webpage that will use the Video Indexer service to display videos of internal company meetings.

You embed the Player widget and the Cognitive Insights widget into the page. You need to configure the widgets to meet the following requirements:

? Ensure that users can search for keywords.

? Display the names and faces of people in the video.

? Show captions in the video in English (United States).

How should you complete the URL for each widget? 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
en-US	<p>Cognitive Insights Widget</p> <p><code>https://www.videoindexer.ai/embed/insights/<accountId>/<videoId>/?widgets=</code> <input type="text" value="Value"/> <code>controls=</code> <input type="text" value="Value"/></p> <p>Player Widget</p> <p><code>https://www.videoindexer.ai/embed/player/<accountId>/<videoId>/? showcaptions=</code> <input type="text" value="Value"/> <code>captions=</code> <input type="text" value="Value"/></p>
false	
people,keywords	
people,search	
search	
true	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 85

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 87

- (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 89

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.

Features	Answer Area
Custom named entity recognition (NER)	Identify code names for internal product development: <input type="text"/>
Key phrase extraction	Identify messages that include credit card numbers: <input type="text"/>
Language detection	
Named Entity Recognition (NER)	
Personally Identifiable Information (PII) detection	
Sentiment analysis	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Features

- Custom named entity recognition (NER)
- Key phrase extraction
- Language detection
- Named Entity Recognition (NER)
- Personally Identifiable Information (PII) detection
- Sentiment analysis

Answer Area

Identify code names for internal product development: Custom named entity recognition (NER)

Identify messages that include credit card numbers: Personally Identifiable Information (PII) detection

NEW QUESTION 94

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.

The screenshot shows a dialog flow in Microsoft Bot Framework Composer. The flow starts with a 'BeginDialog' event, followed by a 'Bot Asks (Text)' step with the prompt 'What is your name?'. This leads to a 'User Input (Text)' step with the assignment 'user.name = Input(Text)'. The flow then goes to another 'Bot Asks (Number)' step with the prompt 'Hello \$(user.name), how old are you?'. This leads to a 'User Input (Number)' step with the assignment 'user.age = Input(Number)'. The configuration panel on the right is for the 'Text input' step, showing 'Property' as 'user.name', 'Output Format' as 'string', and 'Value' as an expression: 'fx = coalesce(@user.Name.@personName)'. The 'Expected responses' are listed as '#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 95

HOTSPOT - (Topic 3)

Select the answer that correctly completes the sentence.

Answer Area

In an infrastructure as a service (IaaS) instance of Microsoft SQL Server on Azure, you manage the

- elastic pool
- MySQL server
- PostgreSQL server
- virtual machine

that hosts SQL Server.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

In an infrastructure as a service (IaaS) instance of Microsoft SQL Server on Azure, you manage the

- elastic pool
- MySQL server
- PostgreSQL server
- virtual machine

that hosts SQL Server.

NEW QUESTION 97

- (Topic 3)

What is the primary purpose of a data warehouse?

- A. to provide storage for transactional line-of-business (LOB) applications
- B. to provide transformation services between source and target data stores
- C. to provide read only storage of relational and non relational historical data
- D. to provide answers to complex queries that rely on data from multiple sources

Answer: C

NEW QUESTION 98

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 99

- (Topic 3)

You have an app named App1 that uses an Azure Cognitive Services model to identify anomalies in a time series data stream. You need to run App1 in a location that has limited connectivity. The solution must minimize costs. What should you use to host the model?

- A. Azure Kubernetes Services (AKS)
- B. a Kubernetes cluster hosted in an Azure Stack Hub integrated system
- C. Azure Container instances
- D. the Docker Engine

Answer: B

NEW QUESTION 101

- (Topic 3)

You have data saved in the following format.

```
FirstName,LastName,Age,LeisureHobby,SportsHobby
John,Smith,23,Reading,Basketball
Ben,Smith,21,Guitar,Curling
```

Which format was used?

- A. CSV
- B. JSON
- C. HTML
- D. YAML

Answer: A

NEW QUESTION 102

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 UserProfile object in the underlying storage layer.	<input type="radio"/>	<input type="radio"/>
The code will create and maintain the ConversationData object in the underlying storage layer.	<input type="radio"/>	<input type="radio"/>
The UserProfile and ConversationData objects will persist when the Bot Framework runtime terminates.	<input type="radio"/>	<input type="radio"/>

- 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 103

HOTSPOT - (Topic 3)

You develop an application that uses the Face API. You need to add multiple images to a person group.

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

```
Parallel.For(0, PersonCount, async i =>
{
    Guid personId = persons[i].PersonId;
    string personImageDir = $"/path/to/person/{i}/images";
    foreach (string imagePath in Directory.GetFiles(personImageDir, "*.jpg"))
    {
        using (  t = File.OpenRead(imagePath))
        {
            await faceClient.PersonGroupPerson. 
        }
    }
});
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Stream

The File.OpenRead(String) method opens an existing file for reading.

Example: Open the stream and read it back. using (FileStream fs = File.OpenRead(path))

Box 2: AddFaceFromStreamAsync

Step 5 on <https://docs.microsoft.com/en-us/azure/cognitive-services/face/face-api-how-to-topics/how-to-add-faces>

NEW QUESTION 107

- (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 containe
- C. Within each indexer definition, schedule the same runtime execution pattern.
- D. Split the data into multiple blob container
- E. Create an indexer for each containe
- 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 folde
- 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 111

- (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 114

- (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 118

- (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 122

- (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 126

- (Topic 3)

You are building a language model by using a Language Understanding service. You create a new Language Understanding resource. You need to add more contributors. What should you use?

- A. a conditional access policy in Azure Active Directory (Azure AD)
- B. the Access control (1AM) page for the authoring resources in the Azure portal
- C. the Access control (1AM) page for the prediction resources in the Azure portal

Answer: B

Explanation:

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-how-to-collaborate>

NEW QUESTION 131

- (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 136

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 138

HOTSPOT - (Topic 3)

You create a knowledge store for Azure Cognitive Search by using the following JSON.

```

"knowledgeStore": {
  "storageConnectionString": "DefaultEndpointsProtocol=https;AccountName=<Acct Name>;AccountKey=<Acct Key>;",
  "projections": [
    {
      "tables": [
        {
          "tableName": "unrelatedDocument",
          "generatedKeyName": "Documentid",
          "source": "/document/pbiShape"
        },
        {
          "tableName": "unrelatedKeyPhrases",
          "generatedKeyName": "KeyPhraseid",
          "source": "/document/pbiShape/keyPhrases"
        }
      ]
    },
    {
      "storageContainer": "unrelatedocrlayout",
      "source": null,
      "sourceContext": "/document/normalized_images/*/layoutText",
      "inputs": [
        {
          "name": "ocrLayoutText",
          "source": "/document/normalized_images/*/layoutText"
        }
      ]
    }
  ],
  "files": []
}

```

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.

Answer Area

There will be [answer choice].

- no projection groups
- one projection group
- two projection groups
- four projection groups

Images will [answer choice]

- not be saved.
- be saved to a blob container.
- be saved to file storage.
- be saved to an Azure Data lake.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

There will be [answer choice].

- no projection groups
- one projection group
- two projection groups
- four projection groups

Images will [answer choice]

- not be saved.
- be saved to a blob container.
- be saved to file storage.
- be saved to an Azure Data lake.

NEW QUESTION 141

HOTSPOT - (Topic 3)

You need to create a new resource that will be used to perform sentiment analysis and optical character recognition (OCR). The solution must meet the following requirements:

- ? Use a single key and endpoint to access multiple services.
- ? Consolidate billing for future services that you might use.
- ? Support the use of Computer Vision in the future.

How should you complete the HTTP request to create the new resource? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

NEW QUESTION 150

- (Topic 3)

Which statement is an example of Data Manipulation Language (DML)?

- A. Revoke
- B. UPDATE
- C. DROP
- D. CREATE

Answer: B

NEW QUESTION 152

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 153

- (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 155

- (Topic 2)

You are developing the chatbot.

You create the following components:

* A QnA Maker resource

* A chatbot by using the Azure Bot Framework SDK.

You need to integrate the components to meet the chatbot requirements. Which property should you use?

A. QnADialogResponseOptions.CardNoMatchText

B. Qna MakerOptions-ScoreThreshold

C. Qna Maker Op t ions StrickFilters

D. QnaMakerOptions.RankerType

Answer: D

Explanation:

Scenario: When the response confidence score is low, ensure that the chatbot can provide other response options to the customers.

When no good match is found by the ranker, the confidence score of 0.0 or "None" is returned and the default response is "No good match found in the KB". You can override this default response in the bot or application code calling the endpoint. Alternately, you can also set the override response in Azure and this changes the default for all knowledge bases deployed in a particular QnA Maker service.

Choosing Ranker type: By default, QnA Maker searches through questions and answers. If you want to search through questions only, to generate an answer, use the RankerType=QuestionOnly in the POST body of the GenerateAnswer request.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/concepts/best-practices>

NEW QUESTION 156

- (Topic 2)

You are developing the chatbot.

You create the following components:

• A QnA Maker resource

• A chatbot by using the Azure Bot Framework SDK

You need to add an additional component to meet the technical requirements and the chatbot requirements. What should you add?

A. Dispatch

B. chatdown

C. Language Understanding

D. Microsoft Translator

Answer: A

Explanation:

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

If a bot uses multiple LUIS models and QnA Maker knowledge bases (knowledge bases), you can use the Dispatch tool to determine which LUIS model or QnA Maker knowledge base best matches the user input. The dispatch tool does this by creating a single LUIS app to route user input to the correct model.

Reference:

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

NEW QUESTION 161

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 164

- (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 165

- (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 166

- (Topic 2)

You are developing the knowledgebase by using Azure Cognitive Search.
 You need to meet the knowledgebase requirements for searching equivalent terms. What should you include in the solution?

- A. a synonym map
- B. a suggester
- C. a custom analyzer
- D. a built-in key phrase extraction skill

Answer: A

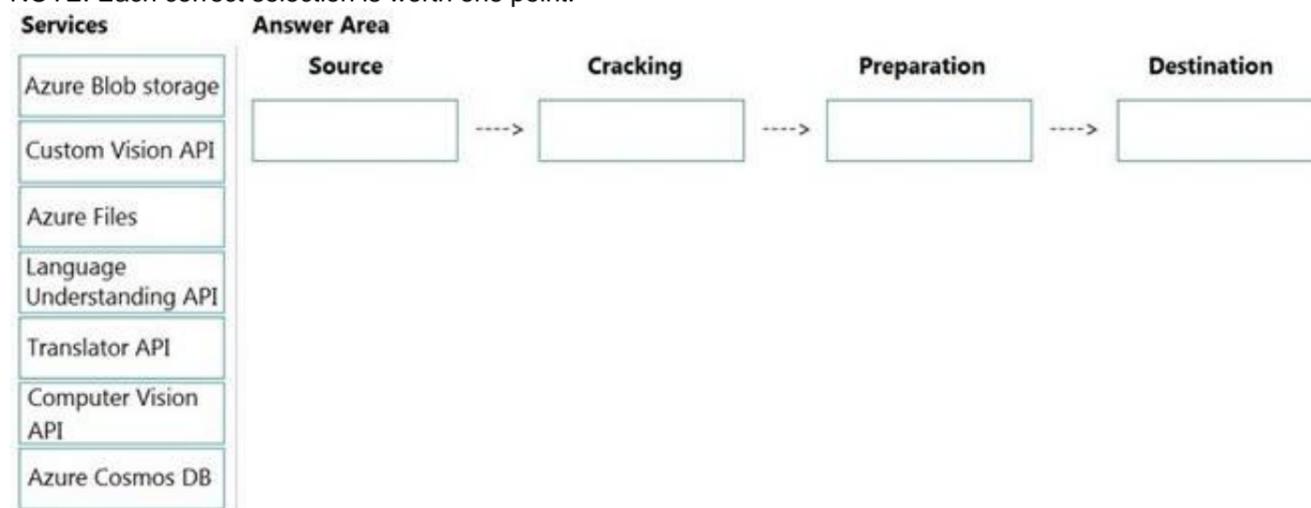
Explanation:

Within a search service, synonym maps are a global resource that associate equivalent terms, expanding the scope of a query without the user having to actually provide the term. For example, assuming "dog", "canine", and "puppy" are mapped synonyms, a query on "canine" will match on a document containing "dog".
 Create synonyms: A synonym map is an asset that can be created once and used by many indexes.
 Reference:
<https://docs.microsoft.com/en-us/azure/search/search-synonyms>

NEW QUESTION 167

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 169

.....

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](#)