

DUMPSQUEEN

Microsoft Azure AI Fundamentals

Microsoft AI-900

Version Demo

Total Demo Questions: 15

Total Premium Questions: 224

Buy Premium PDF

<https://dumpsqueen.com>

support@dumpsqueen.com

dumpsqueen.com

Topic Break Down

Topic	No. of Questions
Topic 2, New Update	94
Topic 3, Describe Artificial Intelligence workloads and considerations	23
Topic 4, Describe fundamental principles of machine learning on Azure	41
Topic 5, Describe features of computer vision workloads on Azure	20
Topic 6, Describe features of Natural Language Processing (NLP) workloads on Azure	27
Topic 7, Describe features of conversational AI workloads on Azure	19
Total	224

QUESTION NO: 1

You need to provide content for a business chatbot that will help answer simple user queries.

What are three ways to create question and answer text by using QnA Maker? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Generate the questions and answers from an existing webpage.
- B. Use automated machine learning to train a model based on a file that contains the questions.
- C. Manually enter the questions and answers.
- D. Connect the bot to the Cortana channel and ask questions by using Cortana.
- E. Import chat content from a predefined data source.

ANSWER: A C E

Explanation:

Automatic extraction

Extract question-answer pairs from semi-structured content, including FAQ pages, support websites, excel files, SharePoint documents, product manuals and policies.

Reference:

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

QUESTION NO: 2 - (HOTSPOT)

HOTSPOT

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
Automated machine learning is the process of automating the time-consuming, iterative tasks of machine learning model development.	<input type="radio"/>	<input type="radio"/>
Automated machine learning can automatically infer the training data from the use case provided.	<input type="radio"/>	<input type="radio"/>
Automated machine learning works by running multiple training iterations that are scored and ranked by the metrics you specify.	<input type="radio"/>	<input type="radio"/>
Automated machine learning enables you to specify a dataset and will automatically understand which label to predict.	<input type="radio"/>	<input type="radio"/>

ANSWER:

Answer Area

Statements	Yes	No
Automated machine learning is the process of automating the time-consuming, iterative tasks of machine learning model development.	<input checked="" type="radio"/>	<input type="radio"/>
Automated machine learning can automatically infer the training data from the use case provided.	<input type="radio"/>	<input checked="" type="radio"/>
Automated machine learning works by running multiple training iterations that are scored and ranked by the metrics you specify.	<input checked="" type="radio"/>	<input type="radio"/>
Automated machine learning enables you to specify a dataset and will automatically understand which label to predict.	<input type="radio"/>	<input checked="" type="radio"/>

Explanation:

Box 1: Yes

Automated machine learning, also referred to as automated ML or AutoML, is the process of automating the time consuming, iterative tasks of machine learning model development. It allows data scientists, analysts, and developers to build ML models with high scale, efficiency, and productivity all while sustaining model quality. Box 2: No

Box 3: Yes

During training, Azure Machine Learning creates a number of pipelines in parallel that try different algorithms and parameters for you. The service iterates through ML algorithms paired with feature selections, where each iteration produces a model with a training score. The higher the score, the better the model is considered to "fit" your data. It will stop once it hits the exit criteria defined in the experiment.

Box 4: No

Apply automated ML when you want Azure Machine Learning to train and tune a model for you using the target metric you specify. The label is the column you want to predict.

Reference: <https://azure.microsoft.com/en-us/services/machine-learning/automatedml/#features>

QUESTION NO: 3 - (DRAG DROP)

Match the Azure Cognitive Services service to the appropriate actions.

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

NOTE: Each correct match is worth one point.

Azure Cognitive Services	Answer Area
Language service	Service Convert spoken requests into text.
Speech	Service Identify the intent of a user's requests.
Translator	Service Apply intent to entities and utterances.

ANSWER:

Azure Cognitive Services	Answer Area
Language service	Translator Convert spoken requests into text.
Speech	Speech Identify the intent of a user's requests.
Translator	Language service Apply intent to entities and utterances.

Explanation:

Azure Cognitive Services	Answer Area
Language service	Translator Convert spoken requests into text.
Speech	Speech Identify the intent of a user's requests.
Translator	Language service Apply intent to entities and utterances.

QUESTION NO: 4 - (HOTSPOT)

HOTSPOT

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements

Yes

No

Providing an explanation of the outcome of a credit loan application is an example of the Microsoft transparency principle for responsible AI.

A triage bot that prioritizes insurance claims based on injuries is an example of the Microsoft reliability and safety principle for responsible AI.

An AI solution that is offered at different prices for different sales territories is an example of the Microsoft inclusiveness principle for responsible AI.

ANSWER:

Answer Area

Statements

Yes

No

Providing an explanation of the outcome of a credit loan application is an example of the Microsoft transparency principle for responsible AI.

A triage bot that prioritizes insurance claims based on injuries is an example of the Microsoft reliability and safety principle for responsible AI.

An AI solution that is offered at different prices for different sales territories is an example of the Microsoft inclusiveness principle for responsible AI.

Explanation:

Box 1: Yes

Achieving transparency helps the team to understand the data and algorithms used to train the model, what transformation logic was applied to the data, the final model generated, and its associated assets. This information offers insights about how the model was created, which allows it to be reproduced in a transparent way.

Box 2: No

A data holder is obligated to protect the data in an AI system, and privacy and security are an integral part of this system. Personal needs to be secured, and it should be accessed in a way that doesn't compromise an individual's privacy.

Box 3: No

Inclusiveness mandates that AI should consider all human races and experiences, and inclusive design practices can help developers to understand and address potential barriers that could unintentionally exclude people. Where possible, speech-to-text, text-to-speech, and visual recognition technology should be used to empower people with hearing, visual, and other impairments.

Reference:

<https://docs.microsoft.com/en-us/azure/cloud-adoption-framework/innovate/best-practices/trusted-ai>

QUESTION NO: 5 - (DRAG DROP)

DRAG DROP

Match the Microsoft guiding principles for responsible AI to the appropriate descriptions.

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

NOTE: Each correct selection is worth one point.

Select and Place:

Principles

Accountability

Fairness

Inclusiveness

Privacy and security

Reliability and safety

Answer Area

Principle

Ensure that AI systems operate as they were originally designed, respond to unanticipated conditions, and resist harmful manipulation.

Principle

Implementing processes to ensure that decisions made by AI systems can be overridden by humans.

Principle

Provide consumers with information and controls over the collection, use, and storage of their data.

ANSWER:

Principles

Accountability

Fairness

Inclusiveness

Privacy and security

Reliability and safety

Answer Area

Reliability and safety

Ensure that AI systems operate as they were originally designed, respond to unanticipated conditions, and resist harmful manipulation.

Accountability

Implementing processes to ensure that decisions made by AI systems can be overridden by humans.

Privacy and security

Provide consumers with information and controls over the collection, use, and storage of their data.

Explanation:

Box 1: Reliability and safety

To build trust, it's critical that AI systems operate reliably, safely, and consistently under normal circumstances and in unexpected conditions. These systems should be able to operate as they were originally designed, respond safely to unanticipated conditions, and resist harmful manipulation.

Box 2: Accountability

The people who design and deploy AI systems must be accountable for how their systems operate. Organizations should draw upon industry standards to develop accountability norms. These norms can ensure that AI systems are not the final authority on any decision that impacts people's lives and that humans maintain meaningful control over otherwise highly autonomous AI systems.

Box 3: Privacy and security

As AI becomes more prevalent, protecting privacy and securing important personal and business information is becoming more critical and complex. With AI, privacy and data security issues require especially close attention because access to data is essential for AI systems to make accurate and informed predictions and decisions about people. AI systems must comply with privacy laws that require transparency about the collection, use, and storage of data and mandate that consumers have appropriate controls to choose how their data is used

Reference:

<https://docs.microsoft.com/en-us/learn/modules/responsible-ai-principles/4-guiding-principles>

QUESTION NO: 6

You are building a knowledge base by using QnA Maker.

Which file format can you use to populate the knowledge base?

- A. PPTX
- B. XML
- C. ZIP
- D. PDF

ANSWER: D

Explanation:

D: Content types of documents you can add to a knowledge base:

Content types include many standard structured documents such as PDF, DOC, and TXT.

Note: The tool supports the following file formats for ingestion:

- .tsv: QnA contained in the format Question(tab)Answer.
- .txt, .docx, .pdf: QnA contained as regular FAQ content--that is, a sequence of questions and answers.

Incorrect Answers:

A: PPTX is the default presentation file format for new PowerPoint presentations. B: It is not possible to ingest xml file directly.

Reference: <https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/concepts/data-sources-and-content>

QUESTION NO: 7

You run a charity event that involves posting photos of people wearing sunglasses on Twitter.

You need to ensure that you only retweet photos that meet the following requirements:

- Include one or more faces.
- Contain at least one person wearing sunglasses.

What should you use to analyze the images?

- A.** the Verify operation in the Face service
- B.** the Detect operation in the Face service
- C.** the Describe Image operation in the Computer Vision service
- D.** the Analyze Image operation in the Computer Vision service

ANSWER: B

Explanation:

Reference:

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

QUESTION NO: 8

In which two scenarios can you use the Form Recognizer service? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A.** Extract the invoice number from an invoice.
- B.** Translate a form from French to English.
- C.** Find image of product in a catalog.
- D.** Identify the retailer from a receipt.

ANSWER: A D

Explanation:

Reference:

<https://azure.microsoft.com/en-gb/services/cognitive-services/form-recognizer/#features>

QUESTION NO: 9

Which two scenarios are examples of a conversational AI workload? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. a telephone answering service that has a pre-recorder message
- B. a chatbot that provides users with the ability to find answers on a website by themselves
- C. telephone voice menus to reduce the load on human resources
- D. a service that creates frequently asked questions (FAQ) documents by crawling public websites

ANSWER: B C

Explanation:

B: A bot is an automated software program designed to perform a particular task. Think of it as a robot without a body.

C: Automated customer interaction is essential to a business of any size. In fact, 61% of consumers prefer to communicate via speech, and most of them prefer self-service. Because customer satisfaction is a priority for all businesses, self-service is a critical facet of any customer-facing communications strategy.

Incorrect Answers:

D: Early bots were comparatively simple, handling repetitive and voluminous tasks with relatively straightforward algorithmic logic. An example would be web crawlers used by search engines to automatically explore and catalog web content.

Reference: <https://docs.microsoft.com/en-us/azure/architecture/data-guide/big-data/ai-overview>
<https://docs.microsoft.com/en-us/azure/architecture/solution-ideas/articles/interactive-voice-response-bot>

QUESTION NO: 10 - (SIMULATION)


To complete the sentence, select the appropriate option in the answer area.

Computer vision capabilities can be Deployed to.....

ANSWER: seetheanswerinbelowExplanation

Explanation:

Integrate a facial recognition feature into an app.



QUESTION NO: 11

Which two scenarios are examples of a conversational AI workload? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. a smart device in the home that responds to questions such as “What will the weather be like today?”
- B. a website that uses a knowledge base to interactively respond to users’ questions
- C. assembly line machinery that autonomously inserts headlamps into cars
- D. monitoring the temperature of machinery to turn on a fan when the temperature reaches a specific threshold

ANSWER: A B

QUESTION NO: 12

You have the process shown in the following exhibit.



Which type of AI solution is shown in the diagram?

- A. a sentiment analysis solution
- B. a chatbot
- C. a machine learning model
- D. a computer vision application

ANSWER: B

QUESTION NO: 13

You are developing a chatbot solution in Azure.

Which service should you use to determine a user's intent?

- A. Translator
- B. QnA Maker
- C. Speech
- D. Language Understanding (LUIS)

ANSWER: D

Explanation:

Language Understanding (LUIS) is a cloud-based API service that applies custom machine-learning intelligence to a user's conversational, natural language text to predict overall meaning, and pull out relevant, detailed information.

Design your LUIS model with categories of user intentions called intents. Each intent needs examples of user utterances. Each utterance can provide data that needs to be extracted with machine-learning entities.

Reference:

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

QUESTION NO: 14

For a machine learning progress, how should you split data for training and evaluation?

- A. Use features for training and labels for evaluation.
- B. Randomly split the data into rows for training and rows for evaluation.
- C. Use labels for training and features for evaluation.
- D. Randomly split the data into columns for training and columns for evaluation.

ANSWER: B

Explanation:

The Split Data module is particularly useful when you need to separate data into training and testing sets. Use the Split Rows option if you want to divide the data into two parts. You can specify the percentage of data to put in each split, but by default, the data is divided 50-50. You can also randomize the selection of rows in each group, and use stratified sampling.

Reference: <https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/split-data>

QUESTION NO: 15

What are three Microsoft guiding principles for responsible AI? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. knowledgeability
- B. decisiveness
- C. inclusiveness
- D. fairness
- E. opinionatedness
- F. reliability and safety

ANSWER: C D F

Explanation:

Reference: <https://docs.microsoft.com/en-us/learn/modules/responsible-ai-principles/4-guiding-principles>