

DUMPS ARENA

Microsoft Azure AI Fundamentals

Microsoft AI-900

Version Demo

Total Demo Questions: 20

Total Premium Questions: 490

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

You need to build an image tagging solution for social media that tags images of your friends automatically.

Which Azure Cognitive Services service should you use?

- A. Face
- B. Form Recognizer
- C. Text Analytics
- D. Computer Vision

ANSWER: A**Explanation:**

To tag photos of your friends automatically, you need a service that can detect faces in an image and then recognize (identify) who those faces belong to. That's exactly what the Azure AI Face service is built for: face detection plus face recognition scenarios like verifying or identifying a person based on their facial features.

Computer Vision can generate general image tags (like "dog", "outdoor", "car"), but it doesn't identify specific people. Form Recognizer is for extracting text/fields from documents, and Text Analytics is for analyzing written language—neither helps with recognizing friends in photos.

So the best fit here is Face. You'd typically detect faces first, then use identification against a known set of people (a person group) to label the image with your friends' names.

References: <https://learn.microsoft.com/en-us/azure/ai-services/computer-vision/overview-identity> and <https://learn.microsoft.com/en-us/azure/ai-services/computer-vision/how-to/identity-detect-faces>

QUESTION NO: 2

A company employs a team of customer service agents to provide telephone and email support to customers. The company develops a webchat bot to provide automated answers to common customer queries.

Which business benefit should the company expect as a result of creating the webchat bot solution?

- A. increased sales
- B. a reduced workload for the customer service agents
- C. improved product reliability

ANSWER: B**Explanation:**

A webchat bot is great at handling the “repeat” questions—things like order status, password resets, business hours, or basic troubleshooting. When the bot answers those common queries automatically, the human agents don’t have to spend as much time on simple, routine requests.

That usually means the support team’s workload drops, and agents can focus on the tougher, more complex customer issues that actually need a person. It can also shorten wait times and improve consistency, but the most direct and expected business benefit here is reducing the amount of work the agents need to do.

Increased sales might happen indirectly, but it’s not the most reliable or primary outcome of adding a support bot. And a chatbot doesn’t really make the product itself more reliable—it just helps customers get answers faster.

References: <https://learn.microsoft.com/en-us/azure/ai-services/bot-service/bot-service-overview> and <https://learn.microsoft.com/en-us/azure/architecture/example-scenario/ai/virtual-agent>

QUESTION NO: 3 - (HOTSPOT)

HOTSPOT

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

Hot Area:

When developing an AI system for self-driving cars, the Microsoft for responsible AI should be applied to ensure consistent operation system during unexpected circumstances.

inclusiveness
accountability
reliability and safety
fairness

principle of the

ANSWER:

When developing an AI system for self-driving cars, the Microsoft for responsible AI should be applied to ensure consistent operation system during unexpected circumstances.

inclusiveness
accountability
reliability and safety
fairness

principle of the

Explanation:

The correct choice is **reliability and safety**. The key clue in the sentence is “ensure consistent operation of the system during unexpected circumstances.” In Microsoft’s Responsible AI principles, reliability and safety is specifically about making sure AI systems behave as intended not only in normal conditions, but also when something unusual happens (edge cases, sensor failures, unexpected inputs, changing environments, etc.). For a self-driving car scenario, this principle is especially important because the system must remain dependable and respond safely when conditions deviate from what it saw during training—like sudden weather changes, unusual road layouts, or unexpected pedestrian behavior.

The other options don’t fit as well: **fairness** focuses on avoiding biased outcomes across groups; **inclusiveness** is about designing systems that empower everyone (including people with disabilities); and **accountability** is about governance and

human oversight for AI decisions. Those are important, but they don't directly target "consistent operation during unexpected circumstances" the way reliability and safety does.

Microsoft's Responsible AI overview and principle descriptions explicitly connect reliability/safety to robust performance and safe behavior under unexpected conditions. See: [Responsible AI overview](#) and [Responsible AI principles](#).

QUESTION NO: 4

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:**

Form Recognizer (now part of Azure AI Document Intelligence) is built for pulling structured data out of documents like invoices, receipts, and forms. So if you need something like the invoice number, that's exactly what it's good at: it reads the document and extracts key fields as data you can store or process.

It also works well with receipts. For example, identifying the retailer (merchant name) from a receipt is a common prebuilt receipt scenario—Form Recognizer can extract merchant/retailer details along with totals, dates, and line items.

On the other hand, translating French to English is a language task (Azure AI Translator), not a document extraction task. And "find image of product in a catalog" is more of a computer vision/search problem (like Azure AI Vision), not something Form Recognizer is designed to do.

References: <https://learn.microsoft.com/en-us/azure/ai-services/document-intelligence/overview> and <https://learn.microsoft.com/en-us/azure/ai-services/document-intelligence/prebuilt/receipt>

QUESTION NO: 5 - (SIMULATION)**SIMULATION**

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

Using Recency, Frequency, and Monetary (RFM) values to identify segments of a customer base is an example of

ANSWER: Seethebelowinexplanation:**Explanation:**

Using Recency, Frequency, and Monetary (RFM) values to identify segments of a customer base is an example of classification.

QUESTION NO: 6

You are developing a Chabot solution in Azure.

Which service should you use to determine a users intent?

- A. Translator
- B. Azure Cognitive Search
- C. Speech
- D. Language

ANSWER: D**Explanation:**

To figure out a user's intent in a chatbot, you need a natural language understanding service—something that can read what the person typed and map it to an “intent” like *BookFlight*, *CheckOrderStatus*, or *ResetPassword*. In Azure, that's done with the Language service (specifically Conversational Language Understanding, which is the newer path that replaced the older LUIS branding).

Translator only converts text between languages, so it doesn't tell you what the user wants. Azure Cognitive Search is for searching and indexing content, not interpreting intent. Speech is about turning audio into text (or text to speech); it can help you capture what the user said, but you still need Language to understand the meaning.

References: <https://learn.microsoft.com/en-us/azure/ai-services/language-service/conversational-language-understanding/overview> and <https://learn.microsoft.com/en-us/azure/ai-services/language-service/overview>

QUESTION NO: 7

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 Azure AI Language Service's question answering? Each correct answer presents a complete solution.

NOTE: Each correct and ask questions by selection is worth one point.

- A. Connect the bot to the Cortana channel using Cortana.
- B. Import chit-chat content from a predefined data source.
- C. Manually enter the questions and answers.
- D. Use Azure Machine Learning Automated ML to train a model based on a file that contains question and answer pairs.
- E. Generate the questions and answers from an existing webpage.

ANSWER: B C E**Explanation:**

In Azure AI Language's Question Answering, you can build your knowledge base content a few straightforward ways. One option is to **manually type** the question-and-answer pairs yourself, which is handy when you already know the exact FAQs you want your bot to handle.

You can also **generate Q&A pairs from an existing webpage**. This is great when your company already has help pages or documentation online—Question Answering can pull in that content and turn it into something the bot can use.

Finally, you can **import predefined chit-chat** content. That gives your bot ready-made small talk (like greetings and basic conversational replies) without you writing it all from scratch. Options like connecting to Cortana or training with Automated ML aren't methods for creating Q&A text inside the Question Answering feature.

References: <https://learn.microsoft.com/en-us/azure/ai-services/language-service/question-answering/overview> and <https://learn.microsoft.com/en-us/azure/ai-services/language-service/question-answering/how-to/chit-chat>

QUESTION NO: 8 - (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.

Statements	Yes	No
You can use QnA Maker to query an Azure SQL database.	<input type="radio"/>	<input type="radio"/>
You should use QnA Maker when you want a knowledge base to provide the same answer to different users who submit similar questions.	<input type="radio"/>	<input type="radio"/>
The QnA Maker service can determine the intent of a user utterance.	<input type="radio"/>	<input type="radio"/>

ANSWER:

Statements	Yes	No
You can use QnA Maker to query an Azure SQL database.	<input type="radio"/>	<input checked="" type="radio"/>
You should use QnA Maker when you want a knowledge base to provide the same answer to different users who submit similar questions.	<input checked="" type="radio"/>	<input type="radio"/>
The QnA Maker service can determine the intent of a user utterance.	<input type="radio"/>	<input checked="" type="radio"/>

Explanation:

QnA Maker (which has since evolved into the Question Answering capability in Azure AI Language) is built for FAQ-style conversational experiences: you create a knowledge base from sources like URLs, documents, or manually authored question-and-answer pairs, and the service uses NLP to match a user's question to the closest Q&A in that knowledge base. It is not a data-access tool and it doesn't execute SQL statements against Azure SQL Database. So the statement about using QnA Maker to query an Azure SQL database is **false** (select **No**).

The second statement describes the core use case for QnA Maker: providing consistent answers to many users who ask the same thing in different ways (for example, "What are your hours?" vs. "When are you open?"). That is exactly what a QnA knowledge base is for, so this statement is **true** (select **Yes**). Microsoft's documentation for QnA/Question Answering describes building a knowledge base from FAQ content and returning the best answer to user questions: [Question Answering overview](#).

The third statement is about identifying a user's *intent*. Intent classification (for example, deciding whether the user wants to "BookFlight" or "CheckWeather") is handled by language understanding/NLU services such as LUIS (historically) or other conversational language understanding capabilities—not by QnA Maker. QnA Maker's job is to retrieve an answer from a curated knowledge base, not to classify intents across multiple actions. Therefore, the statement that QnA Maker can determine intent is **false** (select **No**). For guidance on choosing between QnA-style answering and intent-based language understanding, see: [Language Understanding \(LUIS\) overview](#) (conceptually relevant even though LUIS is being retired/migrated).

QUESTION NO: 9

What are two metrics that you can use to evaluate a regression model? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. coefficient of determination (R2)
- B. F1 score
- C. root mean squared error (RMSE)
- D. area under curve (AUC)
- E. balanced accuracy

ANSWER: A C

Explanation:

For regression, you're predicting a number (like price, temperature, or demand), so you want metrics that tell you how close your predicted numbers are to the real ones. **R-squared (R²)** is a handy "fit" score: it describes how much of the variation in the target value your model can explain. Higher is better, and values close to 1 mean the model is doing a great job capturing the pattern.

RMSE (Root Mean Squared Error) is more direct: it measures the typical size of your prediction errors. It squares the errors first, so big mistakes get penalized more, which is useful when large errors are especially bad. Lower RMSE means your predictions are, on average, closer to the true values.

The other options (F1, AUC, balanced accuracy) are mainly for classification problems where you're predicting categories, not continuous numbers.

References: <https://learn.microsoft.com/en-us/dotnet/machine-learning/resources/metrics>, <https://learn.microsoft.com/en-us/azure/machine-learning/how-to-understand-regression-metrics>

QUESTION NO: 10

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:**

In QnA Maker (part of Azure Cognitive Service for Language), you can build your Q&A knowledge base a few different ways. One common option is to point it at an existing FAQ or support webpage, and it will automatically extract question-and-answer pairs from that content. That's exactly what option A describes.

You can also type your own Q&A pairs directly in the portal. This is useful when you have internal knowledge or very specific answers that aren't already written down somewhere public. That matches option C.

Finally, QnA Maker includes a built-in "chit-chat" feature where you can import a predefined set of small-talk Q&A (like greetings and polite responses). It's a quick way to make the bot feel more natural without writing everything yourself, which is option E.

The other choices don't fit: QnA Maker doesn't rely on Automated ML training from a questions-only file (B), and connecting to Cortana (D) is about channels, not creating Q&A content.

References: <https://learn.microsoft.com/en-us/azure/ai-services/language-service/question-answering/concepts/content-types> and <https://learn.microsoft.com/en-us/azure/ai-services/language-service/question-answering/how-to/chit-chat>

QUESTION NO: 11

You plan to develop a bot that will enable users to query a knowledge base by using natural language processing.

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

NOTE: Each correct selection is worth one point.

- A. Language Service
- B. Azure Bot Service
- C. Form Recognizer
- D. Anomaly Detector

ANSWER: A B**Explanation:**

To build a bot that lets people ask questions in plain language and get answers from a knowledge base, you typically need two pieces: something to host/run the bot and something to understand the user's language (and connect it to your Q&A content).

Azure Bot Service is the core service for creating and deploying the bot itself, handling channels like Teams or web chat, and managing the bot runtime.

Language Service (Azure AI Language) covers natural language features like question answering and intent understanding, which is what you'd use to interpret the user's message and query a knowledge base.

Options like **Form Recognizer** (document extraction) and **Anomaly Detector** (spotting unusual patterns in time-series data) don't help with chatting over a knowledge base, so they're not a fit here.

References: <https://learn.microsoft.com/en-us/azure/bot-service/bot-service-overview-introduction?view=azure-bot-service-4.0> and <https://learn.microsoft.com/en-us/azure/ai-services/language-service/question-answering/overview>

QUESTION NO: 12

You are building an AI-based app.

You need to ensure that the app uses the principles for responsible AI.

Which two principles should you follow? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Implement an Agile software development methodology
- B. Implement a process of AI model validation as part of the software review process

- C. Establish a risk governance committee that includes members of the legal team, members of the risk management team, and a privacy officer
- D. Prevent the disclosure of the use of AI-based algorithms for automated decision making

ANSWER: B C

Explanation:

For responsible AI, you don't just "build the model and ship it." You need checks that help you catch issues like bias, unreliable predictions, privacy problems, and unexpected behavior before (and after) release. That's why having a formal model validation step as part of your review process is a solid fit—it supports reliability and safety, and helps you document what the model does well (and where it doesn't).

It also helps a lot to have clear governance around AI risk. A committee that includes legal, risk, and privacy roles is a practical way to make sure you're thinking about compliance, data protection, transparency expectations, and how to handle real-world impacts. This aligns well with responsible AI ideas like accountability and privacy/security.

Agile can be useful, but it's not a responsible AI principle by itself. And hiding the fact that you're using AI for automated decisions is basically the opposite of transparency.

References: <https://learn.microsoft.com/en-us/training/modules/responsible-ai-principles/> and <https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/innovate/best-practices/trusted-ai>

QUESTION NO: 13

You need to reduce the load on telephone operators by implementing a chatbot to answer simple questions with predefined answers.

Which two AI service should you use to achieve the goal? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Text Analytics
- B. QnA Maker
- C. Azure Bot Service
- D. Translator

ANSWER: B C

Explanation:

For a chatbot that answers common, repeatable questions with pre-written responses, you typically need two pieces: something to store and match the FAQ-style questions, and something to host the actual bot experience.

QnA Maker (now part of Azure AI Language "question answering") is built for exactly this. You load in your FAQs or a knowledge base, and it returns the best matching answer when a user asks a question in their own words.

Azure Bot Service is what turns that Q&A capability into a working chatbot that can be deployed to channels like a website chat widget, Teams, or other messaging apps. It handles the conversation plumbing and connects to your QnA knowledge base.

Text Analytics and Translator are useful in other scenarios (sentiment, language detection, translation), but they don't by themselves provide a predefined-answer FAQ bot.

References: <https://learn.microsoft.com/en-us/azure/ai-services/language-service/question-answering/overview> and <https://learn.microsoft.com/en-us/azure/bot-service/bot-service-overview>

QUESTION NO: 14 - (DRAG DROP)

DRAG DROP

Match the facial recognition tasks to the appropriate questions.

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

NOTE: Each correct selection is worth one point.

Select and Place:

Tasks

- grouping
- identification
- similarity
- verification

Answer Area

- Task Do two images of a face belong to the same person?
- Task Does this person look like other people?
- Task Do all the faces belong together?
- Task Who is this person in this group of people?

ANSWER:

Tasks

grouping
identification
similarity
verification

Answer Area

verification
similarity
grouping
identification

Do two images of a face belong to the same person?

Does this person look like other people?

Do all the faces belong together?

Who is this person in this group of people?

Explanation:

The correct way to think about these Face API tasks is by the kind of comparison you're trying to do: 1:1, 1:many, "find similar," or "cluster." For the question "Do two images of a face belong to the same person?", that's a classic 1:1 scenario, which maps to **verification**. Verification doesn't try to name the person; it just returns a confidence score indicating whether the two faces are likely the same individual.

For "Does this person look like other people?", you're not asking "is it the same person," you're asking "who else looks similar." That maps to **similarity**, which is used to find faces that resemble a given face (for example, returning a ranked list of similar faces from a set).

For "Do all the faces belong together?", you're asking the service to automatically cluster faces into groups where each group represents the same person, without providing names. That is **grouping**. Grouping is useful when you have a pile of photos and want them organized by person, but you don't necessarily know who the people are.

Finally, "Who is this person in this group of people?" implies you have a known set of people (a person group) and want to match an unknown face to one of those known identities. That is **identification** (a 1:many match against a trained group). Microsoft documents these capabilities under Face features such as verify, identify, find similar, and group. See: [Face recognition concepts \(Microsoft Learn\)](#) and [Detect, analyze, and identify faces \(Microsoft Learn\)](#).

QUESTION NO: 15

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:

You'd use the **Detect** operation in the Azure Face service. Detect is designed to find faces in an image and return details about each face it finds. That directly covers the first requirement ("include one or more faces").

It also helps with the sunglasses requirement because face detection can return face attributes (when requested), including whether the person is wearing glasses and the type (for example, sunglasses). So you can programmatically check: "Did Detect find at least one face?" and "Does any detected face have glasses = sunglasses?" and then decide whether to retweet.

The **Verify** operation isn't the right tool here because it's meant for comparing two faces (or a face to a person) to see if they're the same individual. Computer Vision's **Describe Image** and **Analyze Image** are great for general scene understanding, but they're not the most reliable or direct way to confirm sunglasses on a face compared to Face detection with attributes.

References: <https://learn.microsoft.com/en-us/azure/ai-services/computer-vision/overview> and <https://learn.microsoft.com/en-us/azure/ai-services/face/overview>

QUESTION NO: 16 - (DRAG DROP)

DRAG DROP

Match the types of machine learning to the appropriate scenarios.

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

NOTE: Each correct selection is worth one point.

Select and Place:

Machine Learning Types	Answer Area
Facial detection	Machine Learning Type Separate images of polar bears and brown bears.
Facial recognition	Machine Learning Type Determine the location of a bear in a photo.
Image classification	Machine Learning Type Determine which pixels in an image are part of a bear.
Object detection	
Optical character recognition (OCR)	
Semantic segmentation	

ANSWER:

Machine Learning Types

Facial detection

Facial recognition

Image classification

Object detection

Optical character recognition (OCR)

Semantic segmentation

Answer Area

Image classification

Separate images of polar bears and brown bears.

Facial detection

Determine the location of a bear in a photo.

Semantic segmentation

Determine which pixels in an image are part of a bear.

Explanation:

The correct matches come down to what kind of “output” you need from the model: a single label for the whole image, a location for an object, or a per-pixel mask.

Separate images of polar bears and brown bears → **Image classification**. Image classification assigns one label to an entire image (for example, “polar bear” vs. “brown bear”). There’s no requirement to draw boxes or identify exactly where the bear is—just decide which class the image belongs to. This aligns with how Azure describes classification as predicting a category for an image. See [Azure AI Vision overview](#).

Determine the location of a bear in a photo → **Object detection**. Object detection is used when you need to find where an object appears in an image, typically returning bounding boxes (and often labels) for each detected instance. “Location” is the key word here: classification alone can’t tell you where the bear is, only that it’s present. Microsoft’s computer vision documentation commonly distinguishes detection as identifying objects and their coordinates. See [Object detection concepts](#).

Determine which pixels in an image are part of a bear → **Semantic segmentation**. Semantic segmentation labels every pixel, producing a detailed mask of the bear region rather than a coarse rectangle. If the question asks which pixels belong to the bear, that’s exactly segmentation’s purpose. This is the most fine-grained of the three tasks. See [Segmentation concepts](#).

Because each scenario’s required output matches the selected technique, the provided drag-and-drop answers are correct.

QUESTION NO: 17

You need to predict the population size of a specific species of animal in an area. Which Azure Machine Learning type should you use?

- A. clustering
- B. regression
- C. classification

ANSWER: B**Explanation:**

Because you're trying to predict a population *size*, you're predicting a numeric value (like 120 animals, 3,500 animals, etc.). That's exactly what **regression** is for in Azure Machine Learning: estimating a continuous number based on input data (for example, habitat size, food availability, weather patterns, and past counts).

Clustering wouldn't fit because it's about grouping similar items when you don't already have labels (like grouping regions with similar animal activity). Classification also isn't the best match because it predicts categories (like "high/medium/low population"), not an actual numeric count.

So for predicting the actual population number, regression is the clean, standard choice. References: <https://learn.microsoft.com/en-us/training/modules/fundamentals-of-machine-learning/2-types-of-machine-learning> and <https://learn.microsoft.com/en-us/azure/machine-learning/concept-ml-algorithms>

QUESTION NO: 18

You need to implement a pre-built solution that will identify well-known brands in digital photographs. Which Azure AI service should you use?

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

ANSWER: C**Explanation:**

For identifying well-known brands (like logos) in photos using a ready-made feature, the best fit is the Azure AI Vision (Computer Vision) service. It has built-in "Brand detection" as part of its image analysis capabilities, so you don't need to collect training images or build your own model first.

Custom Vision is more for when you want to train a custom image classifier or object detector (for example, your own product packaging or a niche logo that isn't commonly recognized). Face is only for detecting and analyzing faces, and Form Recognizer (now Azure AI Document Intelligence) is aimed at extracting text and structure from documents like invoices and receipts—not brands in photos.

Reference: <https://learn.microsoft.com/en-us/azure/ai-services/computer-vision/overview>

Brand detection details: <https://learn.microsoft.com/en-us/azure/ai-services/computer-vision/concept-brand-detection>

QUESTION NO: 19

Which AI service can you use to interpret the meaning of a user input such as "Call me back later?"

- A. Translator
- B. Text Analytics
- C. Speech

D. Language Understanding (LUIS)**ANSWER: D****Explanation:**

To interpret what someone *means* by a phrase like “Call me back later?”, you need a service that can recognize intent from natural language. That’s exactly what Language Understanding (LUIS) was built for: you train it with example phrases, and it learns to map new user messages to an intent (like “request_callback”) and extract useful details (entities) if needed.

The other options don’t really fit. Translator is for converting text between languages, Speech is mainly about turning audio into text (or text into audio), and Text Analytics is great for things like sentiment, key phrases, and entity recognition—but it doesn’t focus on identifying the user’s intent in a conversational way.

So if your goal is understanding the meaning behind the user’s message so your app can react properly, LUIS is the right choice.

Reference: <https://learn.microsoft.com/en-us/azure/ai-services/luis/what-is-luis>

QUESTION NO: 20

You are processing photos of runners in a race.

You need to read the numbers on the runners’ shirts to identity the runners in the photos.

Which type of computer vision should you use?

- A. facial recognition
- B. optical character recognition (OCR)
- C. semantic segmentation
- D. object detection

ANSWER: B**Explanation:**

To read the numbers printed on the runners’ shirts, you need a feature that can find and extract text from an image. That’s exactly what Optical Character Recognition (OCR) is built for—it detects text in a photo and converts it into readable characters you can store and match to runner IDs.

The other options don’t really fit. Facial recognition is about identifying people by their faces, not reading bib numbers. Object detection can tell you “there’s a person” or even “there’s a shirt,” but it won’t reliably turn the printed number into actual digits. Semantic segmentation focuses on labeling pixels (like separating “runner” from “background”), which still doesn’t solve the “read the number” part.

In Azure, you’d typically use the Computer Vision service (or Azure AI Vision) OCR/Read capability for this kind of task.

Reference: <https://learn.microsoft.com/en-us/azure/ai-services/computer-vision/overview-ocr>