



Designing and Implementing an Azure AI Solution

ASSESSMENT GUIDE



Overview

This document provides information and guidance on how to develop formative and summative assessments for AI-102: Develop AI solutions in Azure. Guidance for assessing student mastery as they progress through the course. This guide includes a set of items for each course topic that can be used throughout the course to monitor student progress and inform your instruction; the assessment items are a mixture of multiple-choice questions and open-ended questions.

This guide is intended to be a reference and starting point for instructors as you plan how to assess your students. As you read through the guide, you may choose to tailor the assessment strategies, including the assessment items and rubric, for your classroom.

Note: These questions are still useful, but use your best judgement in how well they align with the current curriculum. For the most up-to-date questions, see the associated module on Learn.

Contents

Overview.....	1
Topic Questions.....	3
Introduction.....	3
Overview of multiple-choice questions.....	3
Overview of open-ended questions	3
1: Introduction to AI and Azure AI Services.....	4
Multiple choice questions.....	4
Open-ended questions	7
2: Develop generative AI apps and agents	9
Multiple choice questions.....	9
Open-ended questions	12
3: Develop natural language solutions in Azure.....	14
Multiple choice questions.....	14
Open-ended questions	23
4: Develop computer vision solutions in Azure.....	27

Multiple choice questions.....	27
Open-ended questions	31
<i>Custom vision</i>	Error! Bookmark not defined.
5: Develop AI information extraction solutions.....	33
Multiple choice questions.....	34
Open-ended questions	35

Topic Questions

Introduction

This section includes multiple choice and open-ended questions that are aligned to the course topics for AI-102: Develop AI solutions in Azure. Each topic includes a sample of multiple-choice questions and open-ended questions.

You are free to use the questions as they are currently presented or modify as appropriate for your classes. The questions do not appear in any other course materials and are designed to supplement the formative assessment opportunities that are integrated directly into Microsoft Learn and the Microsoft Official Course (e.g. Knowledge Checks, “Try-It” activities, Exercises, Walkthroughs, Labs). They are also designed to allow for easy integration into an online quiz through [Microsoft Forms](#) or through your institution’s Learning Management System (LMS). Each set of topic questions represents indicative coverage for each topic objective domain.

Assigning the topic questions to students as an independent activity will enable you to collect data about individual student progress. However, we recommend that you set aside class time to review answers and address any common student misconceptions, as later topics depend on knowledge and understanding gained earlier in the course.

Overview of multiple-choice questions

The multiple-choice questions require one or more answer responses and will include plausible distractors. These are set at a level slightly lower than the multiple-choice questions in the [AI-102 certification exam](#).

Overview of open-ended questions

The open-ended questions present further challenges beyond single answer responses and include scenario-based questions. These questions give students the opportunity to demonstrate critical thinking through their responses. As the Capstone project for the course is primarily a cloud solution design activity, the open-ended questions at the end of each topic are a useful approach to help build student confidence and critical thinking to encompass a range of potential solutions.

For their responses, students should be encouraged to explore multiple cloud service solutions and adopt a design-first approach before settling on a potential solution. Exploring the official [Microsoft Azure documentation](#) is a great place for students to research and investigate the different cloud services that are available.

1: Introduction to AI and Azure AI Services

Multiple choice questions

1. Which of the capabilities does AI typically exhibit?

Select all that apply.

- a. Visual perception (correct answer)**
- b. Emotional Analysis
- c. Decision making (correct answer)**
- d. Text or Speech Analysis (correct answer)**

2. Data science, machine learning, and artificial intelligence relate to each other in what way?

Select the correct option.

- a. Machine learning is built on artificial intelligence. Artificial intelligence is a subset of data science
- b. Data science is built on machine learning. Machine learning is a subset of artificial intelligence
- c. Artificial intelligence is built on machine learning. Machine learning is a subset of data science (correct answer)**
- d. Data science is built on artificial intelligence. Artificial intelligence is a subset of machine learning

3. To become an AI Engineer, you must first become an expert in data science and machine learning.

Select the correct option.

- a. False. With prepackaged software services that encapsulate AI capabilities, you can integrate AI solutions with a conceptual understanding of data science and machine learning (correct answer)**
- b. False. You must be an expert in machine learning before you can implement an AI solution, but you do not need to be an expert in data science
- c. True. To create an integrated AI solution, you must start from raw data, create machine learning models, and design AI features to become an AI engineer
- d. True. While prepacked software services that encapsulate AI capabilities exist, you cannot integrate them into AI solutions unless you have expert knowledge in data science and machine learning

4. Which of the following features does Azure Machine Learning provide?

Select all that apply.

- a. Automated machine learning (correct answer)**

- b. **Azure machine learning designer (correct answer)**
- c. **Data and compute management (correct answer)**
- d. **Pipelines (correct answer)**

5. Which AI service would you use to generate new content?

Select the correct option.

- a. **Azure OpenAI (correct answer)**
- b. Azure AI Search
- c. Azure AI Vision
- d. Azure AI Content Safety

6. Which of the following best illustrates the Fairness principle in Responsible AI?

Select the correct option.

- a. **An AI system for loan approval that uses a person's credit score and financial history, but not gender or ethnicity, to predict creditworthiness (correct answer)**
- b. A machine learning model for a self-driving car that prioritizes safety by undergoing extensive testing before deployment
- c. An AI-based healthcare application that uses encrypted patient data to maintain privacy and security.
- d. A virtual assistant that uses language understandable to users of all ages to ensure inclusiveness

7. Which of the following Azure AI services provides an AI-enriched indexing for search and knowledge mining?

Select the correct option.

- a. Azure Immersive Reader
- b. **Azure AI Search (correct answer)**
- c. Azure Power Apps
- d. Azure AI Content Safety

Getting started with AI services

1. After creating an Azure AI Service resource, which two are typically needed to access the resource from a client application?

Select the correct option.

- a. The endpoint and a shared password
- b. The endpoint and a valid subscription key (correct answer)**
- c. A user name and password unique to each user
- d. A valid subscription key and a password

2. What is the key difference between single-service and multi-service resources in Azure AI?

Select the correct option.

- a. Single-service resources offer a unique key and endpoint for each service, whereas multi-service resources use a single key and endpoint for multiple services (correct answer)**
- b. Multi-service resources allow for integrated analytics, while single-service resources focus on data storage
- c. Multi-service resources are only available in the free tier, while single-service resources are not
- d. Single-service resources provide access to external databases, whereas multi-service resources do not

3. Which of the following methods can be used to interact with Azure AI services REST APIs for consuming AI functionalities?

Select the correct option.

- a. Sending data in XML format using FTP
- b. Submitting data in JSON format over an HTTP POST request (correct answer)**
- c. Receiving responses as binary data over SMTP
- d. Uploading data directly through a graphical user interface

4. Which programming languages have an available SDK for Azure AI services?

Select all that apply.

- a. PHP
- b. Microsoft C# (.NET Core) (correct answer)**
- c. Python (correct answer)**
- d. Ruby

Open-ended questions

1. A client wants to allow their users to type or speak in many languages, but store the information into their application in English. Which Azure service(s) should they use?

Answer: (Using Azure AI Services, specifically the Speech to Text, Speech Translation, and Translator services, will yield this outcome.)

2. What are the main differences between data science, machine learning, and artificial intelligence?

Answer: (Data science focuses on the processing and analysis of data to define patterns and models. Machine learning is a subset of data science that deals with the training and validation of predictive models. Artificial intelligence builds on machine learning to create software that emulates one or more characteristic of human intelligence.)

3. Why is it important for software engineers to consider Responsible AI and ethics when designing AI systems?

Answer: (AI systems tend to inform decisions based on the data with which they were trained. Understanding the ethical implications of how an AI system will use that data to inform decisions will have a direct impact on the individuals influenced by those decisions.)

4. How would a client integrate a machine learning model built with Azure Machine Learning into their application?

Answer: (Using Azure Machine Learning, the model can be published as a web service and then called like any other web service from the application.)

5. How can a company ensure that its use of AI in customer service respects the Transparency principle in Responsible AI?

Answer: (A company can uphold the Transparency principle by clearly communicating to customers how the AI functions, what data it uses, and how it affects their interactions.)

Getting started with AI services

1. What are the advantages of creating a multi-service resource in Azure AI Services compared to single-service resources for each AI service?

Answer: (Creating a multi-service resource in Azure AI Services allows for access to various AI services with one key and endpoint and provide consolidated billing for easier cost management.)

2. Describe the steps a developer would take to use a REST API provided by Azure AI services in their application.

Answer: (Should include a statement about how the developer sends an appropriate HTTP request (POST, PUT, or GET) containing JSON-formatted data to the Azure AI service. They would handle authentication, submit the request to the service's endpoint, and parse the JSON in the received HTTP response to integrate the AI service's output into their application.)

3. Explain how an SDK can aid a developer in creating applications with Azure AI services, and give examples of what the SDK might simplify compared to using REST APIs directly.

Answer: (Should include a statement about how an SDK, or Software Development Kit, provides language-specific libraries that abstract the details of the REST interfaces, making it easier for developers to build complex solutions. Thus, instead of manually constructing HTTP requests and handling responses, a developer can use the methods and classes provided by the SDK, which handle these tasks automatically.)

2: Develop generative AI apps and agents

Multiple choice questions

1. Which of the following best describe Generative AI?

Select the correct option.

- a. The field of computer science that seeks to create intelligent machines that can replicate or exceed human intelligence
- b. Subset of AI that enables machines to learn from existing data and improve upon that data to make decisions or predictions
- c. A machine learning technique in which layers of neural networks are used to process data and make decisions
- d. Create new written, visual, and auditory content given prompts or existing data (correct answer)**

2. You are planning to develop a support application that can summarize text and suggest code. Which of the following best describes the benefits and features of using Azure AI Foundry?

Select the correct option.

- a. Azure AI Foundry provides text generation capabilities but lacks integration with other cloud services
- b. The use of Azure AI Foundry requires extensive machine learning experience
- c. Azure AI Foundry brings generative AI models to the Azure platform, providing security, scalability, and integration with other Azure services (correct answer)**
- d. Azure AI Foundry offers generative AI models that are restricted to language processing

3. Which of the following statements best describes the base models available in Azure OpenAI?

Select the correct option.

- a. Azure OpenAI only offers GPT-4 models, which are specialized in generating natural language and code completions
- b. Azure OpenAI includes GPT-4, GPT 3.5, Embeddings, and DALL-E models, each with unique capabilities such as language generation, code completions, text analytics, and image generation (correct answer)**
- c. Azure OpenAI primarily focuses on Embeddings models for text analytics
- d. All models in Azure OpenAI are variants of ChatGPT models

4. Which of the following factors influence the output completion provided by the generative AI model?

[Develop AI solutions in Azure](#)

Select the correct option.

- a. The location of the server on which the model is hosted
- b. The text in the prompt, prompt engineering techniques, and the set parameters for the interaction (correct answer)**
- c. The network speed and the training date for the model
- d. The time of day during which the prompt is submitted

Prompt engineering

1. What is prompt engineering in the context of interacting with language models Foundry?

Select the correct option.

- a. The process of prompt engineering involves designing and optimizing prompts to utilize LLMs effectively, improving the model's performance on specific tasks and mitigating bias (correct answer)**
- b. The process of prompt engineering involves the practice of providing random prompts to see how the AI model responds improving their efficiency
- c. Prompt engineering is the technical term for troubleshooting AI models when they do not understand prompts
- d. Prompt engineering is the design of the algorithms used by the AI models in the processing of the input prompt

3. How does the temperature parameter affect the responses generated by Azure OpenAI models?

Select the correct option.

- a. Lower values of temperature reduce randomness in the response, making it more consistent and focused (correct answer)**
- b. Higher values of temperature restrict the AI to generate only factual information
- c. The temperature parameter controls the length of the response
- d. The temperature parameter has no effect on the responses generated rather the speed of the response

4. What is the purpose of using section markers within prompts for generative models?

Select the correct option.

- a. They help the model differentiate between instructions and content (correct answer)**
- b. They are used to separate different languages within the prompt
- c. Section markers indicate to the model the randomness of the response

d. Section markers have no functional purpose in the prompt rather than make it easier for users to construct prompt

5. What is the role of primary, supporting, and grounding content in prompt engineering?

Select the correct option.

a. Primary content is the main subject of the prompt, supporting content provides additional details that refine the response, and grounding content offers a factual basis from which the model can draw answers (correct answer)

b. Primary content refers to detailed instructions for the model, supporting content adds additional context, and grounding content is irrelevant to the prompt

c. Primary content is the main subject of the prompt, supporting content provides offers a factual basis from which the model can draw answers, and grounding content adds additional context

d. Primary content and supporting content are used to define context for the prompt, while grounding content provides the generative model details

6. What is the function of cues in prompt engineering?

Select the correct option.

a. They are leading words or phrases that guide the model to build upon and shape the response in a desired direction (correct answer)

b. Cues are used to signal the end of a prompt

c. Cues refer to hidden keywords that unlock advanced features of the Azure OpenAI models

d. Cues are error messages returned by the model when a prompt is incorrectly formatted

RAG

1. What are the differences between fine-tuning and using your own data with Azure AI Foundry?

Select the correct option.

a. Fine-tuning involves training an existing model like gpt-35-turbo with additional data for better prompt responses, whereas using your own data with the stateless API does not require custom model training (correct answer)

b. Fine-tuning creates a custom model by training with additional data and can yield higher quality responses with fewer examples, while using your own data involves training a custom model each time

c. Fine-tuning is a quick and cost-effective way to use Azure OpenAI with your own data, whereas using the RAG file connection is more costly and time-intensive

d. Both fine-tuning and using your own data require extensive training with large datasets

2. How can you add your own data to Azure OpenAI to enhance the model's responses in the Chat playground?

Select the correct option.

- a. Through Azure AI Studio, by uploading files, using blob storage, or connecting to an existing Azure AI Search index (correct answer)**
- b. By manually entering data into each prompt before sending it to the model
- c. Adding your own data is not supported; the model can only use pre-trained information
- d. You cannot add your own data through the Chat playground, you have to use either the REST API or the SDK

Open-ended questions

1. Explain what generative AI is and how is it different from AI, machine learning and deep learning.

Answer: (Should include a statement that says that generative AI create new written, visual, and auditory content given prompts or existing data. AI is The field of computer science that seeks to create intelligent machines that can replicate or exceed human intelligence. Machine learning is a subset of AI that enables machines to learn from existing data and improve upon that data to make decisions or predictions. Deep learning is a machine learning technique in which layers of neural networks are used to process data and make decisions.)

2. How do the different models available in Azure OpenAI, such as GPT-4, GPT 3.5, Embeddings, and DALL-E, vary in their application and performance, and what factors should be considered when choosing a model for a specific task.

Answer: (Should include a statement that include the difference between the models, GPT-4 models are the latest generation of generative pretrained (GPT) models that can generate natural language and code completion. GPT 3.5 models can generate natural language and code completions. GPT-35-turbo models are optimized for chat-based interactions and work well in most generative AI scenarios. Embeddings models convert text into numeric vectors and are useful in language analytics scenarios. DALL-E models are used to generate images based on natural language prompts.)

3. How can integrating Azure AI Foundry enhance the capabilities of an application, and what are some potential use cases?

Answer: (Should include a statement that says that the endpoints can be integrated into applications to automate content creation, enhance user interactions with chatbots, and provide code suggestions. Potential use cases include writing assistants, programming help tools, and customer service bots.)

Prompt Engineering

1. Discuss how can you use prompt engineering to provide a more relevant responses from Azure OpenAI models?

Develop AI solutions in Azure

Answer: (Should include a statement that discuss different prompt engineering strategies, like prompt should be clear and specific about what is being asked. Providing sufficient contextual information in the prompt with the use of primary, supporting and grounding content. The use of cues to steer the model in a certain direction. Request output composition. Use of instructional language and leveraging examples.)

2. Discuss how the inclusion of primary, supporting, and grounding content in your prompts can lead to more effective responses from Azure OpenAI models?

Answer: (Should include a statement that says that primary, supporting, and grounding content in prompts is used in prompt engineering to provide more relevant responses. Primary content is the main content to be summarize, translated or processed, provide the main subject of the prompt. Supporting content provide clarity and specification. Grounding content define the scope.)

3. How can cues be effectively used in prompt engineering to enhance the quality of responses from Azure OpenAI models?

Answer: (Should include a statement that says that the use of cues in prompt engineering involves strategically placing leading words or phrases that direct the Azure OpenAI model towards generating a specific type of response. They can be particularly useful when generating code.)

RAG

1. Describe the process of preparing and connecting your own data to Foundry for improved response accuracy.

Answer: (Should include a statement that discusses how to add your data to Azure OpenAI. This includes adding data through Azure AI Studio in the Chat playground by uploading files directly, selecting data from blob storage, or linking to an Azure AI Search index.)

2. What are the key differences between fine-tuning an existing Azure model and using your own data with Azure AI Foundry?

Answer: (Should include a statement that discusses the difference between fine tuning and using your own data as the degree of customization, the time and resources required, and the specific use cases they are best suited for. Fine-tuning involves training an existing model like gpt-35-turbo with additional data for better prompt responses, whereas using your own data with the stateless API does not require custom model training.)

3: Develop natural language solutions in Azure

Multiple choice questions

Analyze text

1. What functionality does the Azure AI Language service provide?

Select all that apply.

- a. Language detection (correct answer)**
- b. Language translation
- c. Sentiment analysis (correct answer)**
- d. Key phrase extraction (correct answer)**

2. A company wants to categorize a collection of documents by language. The documents each represent an email from a customer. Using the detect language feature of Azure AI Language service, will you be able to detect the language of each email separately?

Select the correct option.

- a. Yes, but you have to send each document individually
- b. Yes, a JSON response will contain results for each document individually (correct answer)**
- c. No, you have to specify the language of each document and detect language only confirms the language
- d. No, there isn't a language detection feature in the Azure AI Language service

3. What is an example of a use case for the Azure AI Language service sentiment analysis feature?

Select the correct option.

- a. To monitor open-ended feedback on a service or product and prioritize those with negative feedback to reach out to those customers. (correct answer)**
- b. To identify key phrases from comments on a product page and prioritize highlighting them in future ads
- c. To localize language based on the region of a customer and use specific dialects in chat bots on a company help page
- d. To categorize customer service feedback automatically and create performance reviews without having to review each piece of feedback manually

4. What summarization services does the Azure AI Language Service offer and is it customizable?

[Develop AI solutions in Azure](#)

Select the correct option.

- a. Only extractive summarization, which can be customized with user data to improve accuracy
- b. Only abstractive summarization, which uses the most important sentences from the source document
- c. Both extractive and abstractive summarization, where extractive uses the source's exact words and abstractive captures the main idea; can be customized with user data (correct answer)**
- d. Both extractive and abstractive summarization, however it cannot be customized with user data

5. What is the purpose of the Personal Identifiable Information (PII) detection feature in Azure AI Language Service?

Select the correct option.

- a. To enhance sensitive information and focus on entities like Organization, Location, and Date
- b. To detect and redact sensitive information including entities such as Person, Phone Number, Email, Address, Credit card, and financial account identification (correct answer)**
- c. To detect and store sensitive information related to medical and health records only
- d. To label sensitive information for marketing purposes, specifically targeting entities like Age, Gender, and Education

Translate text

1. What are the three features of the Translator service in Azure AI Services?

Select all that apply.

- a. Translation (correct answer)**
- b. Transcription
- c. Language detection (correct answer)**
- d. Transliteration (correct answer)**

2. Which of the following parameters of the Translator service should you use to avoid displaying translations of profanities?

Select the correct option.

- a. includeAlignment
- b. profanityMarker (correct answer)**
- c. customTranslate
- d. A parameter to avoid profanities does not exist

3. What are the correct steps for creating a custom translation model using the Azure Translator service?

Select the correct option.

- a. Link a workspace to your Azure AI Translator resource, create a project, use the Custom Translator portal, and then train a model after uploading training data files
- b. Use the Custom Translator portal, link a workspace to your Azure AI Translator resource, upload training data files, create a project, and then train a model
- c. Use the Custom Translator portal, link a workspace to your Azure AI Translator resource, create a project, upload training data files, and then train a model (correct answer)**
- d. Create a project, link a workspace to your Azure AI Translator resource, use the Custom Translator portal, upload training data files, and then train a model

4. Which of the following statements accurately describes the Azure AI Translator parameters and their effects?

Select the correct option.

- a. The includeAlignment parameter can be used to show how words in English map to characters in languages like Simplified Chinese, aiding in understanding the translation relationship between source and target texts (correct answer)**
- b. The includeSentenceLength parameter, when set to true, obscures any profanities found in the source text by default
- c. The profanityAction parameter can be set to Deleted to ensure that sentence length information is included in the translation results
- d. Setting the profanityMarker parameter to Tag will replace spaces in languages that do not use them with XML tags to indicate word boundaries

Question answering

1. A knowledge base is what?

Select the correct option.

- a. Question and answer pairs that can be queried using natural language input (correct answer)**
- b. Question and answer pairs that can be queried using exact language input
- c. A list of facts that can be queried using natural language input
- d. A list of questions that can be queried using exact language input

2. What are two ways you can populate a knowledge base with existing question and answer pairs?

Select all that apply.

- a. **Import questions and answers from existing web pages or documents (correct answer)**
- b. Import answers from an existing web pages or documents and use AI to create the questions
- c. **Add pre-defined "chit-chat" pairs that include common conversational questions and responses (correct answer)**
- d. You cannot populate a knowledge base with existing question and answer pairs

3. When should you use question answering over language understanding in Azure AI Language service, and what are the primary differences between the two?

Select the correct option.

- a. **Use question answering when the user submits a query expecting a direct response; it matches questions to answers using natural language understanding. Use language understanding when the user's utterance requires an action or response, as it interprets intent and entities (correct answer)**
- b. Question answering is suitable when the user needs the service to perform an action, while language understanding is designed for when a user submits a question looking for a static answer
- c. Language understanding uses a knowledge base to provide answers to users' questions, whereas question answering interprets natural language to perform actions based on the user's intent
- d. Both question answering and language understanding are designed to provide static answers to known questions, and there is no significant difference in their usage patterns or responses

4. In Azure AI Language service, what action does the Question Answering enable that the Language Understanding service does not?

Select the correct option.

- a. User submits an utterance, expecting an appropriate response or action
- b. Response indicates the most likely intent and referenced entities
- c. **Response is a static answer to a known question (correct answer)**

5. To create a knowledge base for question answering in Azure AI Language Studio, what are the correct sequences of steps to follow?

Select the correct option.

- a. In Language Studio, create your Azure AI Language resource, create a custom question answering project, populate the knowledge base by uploading document files, and then edit question and answer pairs
- b. Create an Azure AI Language service resource, in Language Studio select it and upload document files, then directly edit question and answer pairs to populate the knowledge base

c. Create an Azure AI Language service resource, use Language Studio to create a custom question answering project, add one or more data sources to populate the knowledge base, and then edit question and answer pairs to finalize the knowledge base (correct answer)

d. Start by editing question and answer pairs in Language Studio, then create a Custom question answering project, and finally link it to your Azure AI Language service resource by importing an FAQ web page

Conversational understanding

1. When building a Language Understanding application, what do we call the phrase that the user might enter?

Select the correct option.

- a. Sentence
- b. Snippet
- c. Utterance (correct answer)**
- d. Intent

2. What are the three entity types?

Select the correct option.

- a. Machine learned, blob, AI
- b. List, blob, RegEx
- c. Machine learned, list, prebuilt (correct answer)**
- d. List, array, tree

3. The Language Understanding service can be deployed in which of the following containers?

Select all that apply.

- a. Local Docker host (correct answer)**
- b. Azure Container Instance (correct answer)**
- c. Azure Docker Container
- d. Azure Kubernetes Service (correct answer)**

4. You want to get only the most likely intent for your natural language query. Which parameter should you use?

Select the correct option.

- a. kind

b. topIntent (correct answer)

- c. query
- d. result

5. You want to filter reviews in your app to know how many people like or dislike a business. What feature would you use?

Select the correct option.

a. Sentiment analysis (correct answer)

- b. Key phrase detection
- c. Entity recognition
- d. Conversational language understanding

6. What is the role of a language model in Natural Language Processing (NLP), and how does Azure facilitate NLP in conversational applications?

Select the correct option.

- a. A language model in NLP is used for syntax highlighting in user input, and Azure's Conversational Language Understanding (CLU) service provides a platform for building database queries in conversational apps
- b. In NLP, a language model is crucial for interpreting user input, which is also known as natural language understanding (NLU); Azure offers the Conversational Language Understanding (CLU) service to build NLU components for conversational applications (correct answer)**
- c. A language model in NLP is primarily used for translating user input from one language to another, and Azure's CLU service aids in this translation for conversational applications
- d. The language model in NLP is designed to correct grammatical errors in user input, and Azure's CLU service is a tool for grammar correction in conversational interfaces

7. What is the correct sequence of steps to create a conversational language understanding model in Azure?

Select the correct option.

- a. Deploy a model to a public endpoint, review predictions, train the model with intents and entities, and then test the model with a dataset
- b. Train a model with sample utterances, test the model interactively or using testing dataset, deploy the model to a public endpoint, and review predictions for further iteration (correct answer)**
- c. Review model predictions first, deploy a trained model to a public endpoint, then train and test the model using sample utterances

- d. Train a model using a testing dataset with known labels, deploy to a public endpoint, and then review predictions to iterate on utterances

Custom text

1. How does custom text classification work in Azure AI Language service?

Select the correct option.

- a. Custom text classification only allows for multiple label classification, meaning each text file can be assigned a combination of predefined classes
- b. Custom text classification works by assigning predefined labels to text, but is limited to a single label per file
- c. Custom text classification in Azure AI Language service can either assign a single label per text file or multiple labels (correct answer)**
- d. Custom text classification defines its own classes and automatically assigns multiple labels to text files without developer intervention

2. What are the key differences between single label and multiple label projects in custom text classification?

Select the correct option.

- a. In single label projects, each file can be classified into multiple categories simultaneously, whereas in multiple label projects, files are restricted to one category only
- b. Multiple label projects are used for simple binary classifications, whereas single label projects are designed for more complex classification tasks with multiple categories
- c. Single label projects allow for only one class to be assigned to each file, while multiple label projects can assign various classes to a single file, with differences also present in labeling, model improvement considerations, and API payload structure (correct answer)**
- d. The primary difference between single and multiple label projects is the complexity of the API payload

3. How does the JSON body structure differ between single label and multiple label classification models?

Select the correct option.

- a. Single label classification models require a 'classes' array inside the 'assets' object, while multiple label classification models do not require the 'classes' array
- b. Multiple label classification models use the 'customMultiLabelClassification' value for 'projectKind' and allow multiple 'category' objects within the 'classes' array for a single document, unlike single label classification models (correct answer)**
- c. The JSON body for single label classification models uses 'customMultiLabelClassification' for the 'projectKind' and includes an array of 'documents' with multiple 'category' objects each

d. There is no structural difference in the JSON body between single and multiple label classification models; the only difference is the value of 'projectKind' in the metadata

4. What is the correct sequence of steps in the Azure AI Language project life cycle for custom text classification?

Select the correct option.

- a. Train model, deploy model, define labels, tag data, view model, improve model, and classify text
- b. Define labels, tag data, view model, train model, improve model, deploy model, and classify text
- c. Define labels, tag data, train model, improve model, view model, deploy model, and classify text
- d. Define labels, tag data, train model, view model, improve model, deploy model, and classify text (correct answer)**

5. What are the recommended proportions for splitting your data into training and testing when preparing dataset datasets for training a custom text classification model?

Select the correct option.

- a. 50% for training and 50% for testing
- b. 80% for training and 20% for testing (correct answer)**
- c. 70% for training and 30% for testing
- d. 60% for training and 40% for testing

6. What are the deployment options for a custom text classification model in Azure AI Language?

Select the correct option.

- a. You can deploy one model with one unique name per project
- b. You are limited to deploying models only via the Language Studio
- c. Each project can create multiple models and deployments, each with their own unique name (correct answer)**
- d. You can only deploy models via the REST API

7. What is the difference between Custom Named Entity Recognition (NER) and built-in NER in Azure AI Language?

Select the correct option.

- a. Custom NER identifies generic entities like names or locations, while built-in NER extracts user-defined entities from documents

b. Built-in NER requires a custom model to be trained, whereas Custom NER is ready to use with no additional setup

c. Built-in NER can recognize entities like persons and locations with minimal setup, while Custom NER allows for identifying specific entities not included in the built-in categories (correct answer)

d. Custom NER is an external service that needs to be integrated with Azure AI, unlike built-in NER which is part of the Azure AI Language services

8. What are key considerations for selecting data and refining entities in Custom Named Entity Recognition (NER)?

Select the correct option.

a. Use a narrow dataset focused on a single source, ensure ambiguous entities for complexity, and train with mainly synthetic data

b. Prioritize high-quality, diverse, and accurately distributed real-world data, with entities defined as broad categories like "Contact info" to cover various types

c. Select high-quality data with diversity and accurate real-world distribution, and define entities as distinct and specific types rather than ambiguous or broad categories (correct answer)

d. Focus on a uniform dataset from similar sources for consistency, use mostly hypothetical data, and define entities in general terms for broader application

9. What are some of the limitations of the Azure AI Language service for Custom Named Entity Recognition (NER) projects?

Select all that apply.

a. A minimum of 10 files and a maximum of 100,000 files for training (correct answer)

b. 10 deployment names per project (correct answer)

c. Only 1 storage account per project (correct answer)

d. You can have up to 100 entity types

Speech

1. What information do you need from the Speech resource to use it from a client application through one of the supported SDKs?

Select all that apply.

a. Endpoint (correct answer)

b. Key (correct answer)

c. Subscription ID

- d. Unique Resource ID

2. When using the text-to-speech feature of the Speech service, where do the voices come from?

Select the correct option.

- a. You have to provide the voices by recording a library of words and sounds
- b. Azure provides one voice that everyone can use
- c. Azure provides synthetic and more natural sounding voices for you to choose from (correct answer)**
- d. You can to link a Voice service to the text-to-speech application to use any of Azure's voices

3. What do you need to use to specify the characteristics of speech produced by the Speech SDK?

Select the correct option.

- a. Speech characteristics cannot be specified
- b. Speech characteristics are randomized
- c. Parameters that specify emotion
- d. Speech Synthesis Markup Language (correct answer)**

4. Which setting should you specify if you want to translate speech from an audio file to text?

Select the correct option.

- a. AudioConfig (correct answer)**
- b. SpeechConfig
- c. SpeechTranslationConfig
- d. TranslationRecognizer

Open-ended questions

Analyze text

1. A company wants store written feedback from customers and send the potentially negative feedback to customer service representatives to follow up with those customers. Regardless of the language the feedback was originally written in, they want to send the feedback in English. What services and features should they use?

Answer: (Should include Azure AI Language sentiment analysis and Translate translation)

2. Imagine you are given all of the chat conversation between you and a friend, what services could you use to find the times you coordinated seeing each other in person?

Answer: (Should include the extract entities feature of the Azure AI Language service. This feature includes identifying a person, location, and datetime that might be useful for finding these meetups.)

3. Describe a scenario where it would be imperative to use the extract linked entities feature of the Azure AI Language service over the extract entities feature, and why.

Answer: (Should include a scenario where one word can have two very different meanings, for example Venus the planet and Venus the Greek goddess.)

4. A healthcare research organization is preparing a dataset of patient feedback for a machine learning project aimed at improving patient care. To maintain patient confidentiality, they need to remove all PII from the feedback before analysis. What services and features should they use?

Answer: (Should include Azure AI Language sentiment PII detection feature to scan and redact sensitive information such as names, phone numbers, email addresses, physical addresses, credit card numbers, and financial account identification from the patient feedback.)

Translate text

1. A company wants to create a Latin script version of feedback that was originally written in Japanese with Japanese script. Does Azure AI Services have support for this? If so, how?

Answer: (Should include a statement that mentions the Transliteration feature of the Translator service in Azure AI Services.)

2. A company has a specific vocabulary that they have for their product and wants to automatically translate their documentation into more than one language. Does Azure AI Services have support for this? If so, how?

Answer: (Should include a statement that mentions the Custom Translator feature in the Translator service in Azure AI Services.)

3. A company is looking to improve their global customer support by creating a custom translation model that understands and translates technical language specific to their products. Does Azure AI Services have support for this? If so, how?

Answer: (Should include a statement that mentions the Custom Translator feature in the Translator service in Azure AI Services. Should explain steps to create a custom translation model)

Question answering

1. What are the primary differences between Question Answering and Language Understanding?

Answer: (Should include a statement that mentions usage patterns (questions vs utterances), query processing (matching questions compared to matching intent), response (static answers vs intent and reference entities), client logic (presents the answer vs performing appropriate action).)

2. What is the difference between Implicit and Explicit feedback?

Answer: (Should include a statement that question answering uses implicit feedback by identifying user-provided questions that have multiple, similarly scored matches in the knowledge base. These are then automatically clustered as alternate phrase suggestions for the possible answers that you can view in the portal and choose to accept or reject. Explicit feedback controls the number of possible question matches returned for the user's input by specifying the top parameter.)

3. When you build a question answering solution, can you train and test a natural language model before publishing it for use in an application or bot?

Answer: (Should include a statement that you can test your knowledge base interactively in the Language Studio portal, submitting questions and reviewing the answers that are returned. You can inspect the results to view their confidence scores as well as other potential answers.)

4. Can you create a question answering solution that handle multi-turn conversations?

Answer: (Should include a statement that yes, you can enable multi-turn responses when importing questions and answers from an existing web page or document based on its structure, or you can explicitly define follow-up prompts and responses for existing question and answer pairs in the Language Studio portal.)

5. Discuss how can you improve the question answering performance in Azure AI Language Service?

Answer: (Should include a statement discuss enabling Active Learning within the Azure AI Language Service suggest alternatives when multiple questions have similar scores for user input, explain difference between implicit and explicit feedback. Discuss creation of Synonyms for terms with the same meaning.)

Conversational understanding

1. When labeling your data to train a model, what are your datasets and how are they split?

Answer: (Should include a statement that mentions testing and training datasets, and the ability to manually or automatically split them.)

2. How can you use patterns to differentiate similar utterances?

Answer: (Should include a statement that mentions separating intent from the entities. For example, a TurnOnDevice intent might have “Turn the {DeviceName} on.” And “Switch the {DeviceName} on.” Where {DeviceName} can be any entity.)

3. Why would you want to deploy a Language Understanding app in a container?

Answer: (Should include a statement that mentions improved security, compliance, diagnostics, or reducing the need to manage your own servers.)

4. What benefit does training a model with the Language Understanding service mean practically speaking?

Answer: (Should include a statement that mentions that it enables you to use customized intents for your model to recognize and use in a client app.)

5. What intent must all models include, and what is it used for?

Answer: (Should include a statement that mentions the *None* intent, and how it’s used for utterances it receives that there is no action for (such as “Hello” or “Goodbye”).

Custom text

1. You are tasked with creating a custom text classification model to categorize customer feedback into various service aspects. Describe the process you would follow using Azure AI Language.

Answer: (Should include a statement that of the life cycle of training a custom text classification model, starting with defining labels, tagging data, training model, reviewing model performance, improving the model, and deploying the model.)

2. A media company wants to classify news articles into different genres automatically. How can they accomplish that using custom text classification.

Answer: (Should include a statement regarding the use of multiple label classification models in custom text classification.)

3. You are developing a custom text classification model for classifying legal documents. What strategies can you use to improve the model’s accuracy and how would these be implemented in Azure AI Language?

Answer: (Should include a statement that mentions reviewing model performance by seeing which classification filed to evaluate to the right label, reviewing label distribution and finding what data to add to improve the performance.)

4. How would a financial institution utilize Custom Named Entity Recognition (NER) in Azure AI Language to enhance the extraction of specific entities from their data?

Answer: (Should include a statement that mention the training a custom NER model detailing the process.)

5. How should a healthcare organization approach data selection and entity refinement when developing a Custom NER model to extract patient information from various health records?

Answer: (Should include a statement that mention the selection of high-quality data with diversity and accurate real-world distribution, and defining entities as distinct and specific types rather than ambiguous or broad categories.)

Speech

1. A company wants to create a speech-based help desk for English and Spanish patrons but only have English-speaking customer service employees. Describe how you would help them with this using Azure AI Services.

Answer: (Should include a statement that mentions speech-to-text translations using the TranslationRecognizer and the SpeechSynthesizer to synthesize an audio stream for each language.)

2. Customers often appreciate when synthesized voices are not only localized to language but also sound realistic and have a similar accent to themselves. How would you accomplish this using the Speech SDK?

Answer: (Should include a statement that mentions the neural voices that are created using deep neural networks to sound more natural. Should also include a statement that mentions the locale and person's name of the neural voice for a more accurate accent.)

3. What feature of the Speech SDK would enable you to create the following spoken dialog between two neural voices and how?

Ariana (cheerfully): "I say tomato:

Guy: "I say tomato (pronounced tom-ah-toe) ... Let's call the whole thing off!"

Answer: (Should include a statement that mentions the Speech Synthesis Markup Language (SSML) and features such as speaking styles, phonemes, and prosody.)

4: Develop computer vision solutions in Azure

Multiple choice questions

Image analysis and Face

1. When analyzing an image with AI Vision, what image features can be retrieved?

Select the correct answer.

- a. **Image description (correct answer)**
- b. **Tags (correct answer)**
- c. **Objects in the image (correct answer)**

- d. A list of all colors
2. An application wants to automatically remove images that may contain inappropriate content from their database. What image feature should they look for in the image analysis response?
Select the correct answer.
- a. This isn't an available feature in AI Vision
 - b. "Objects" and compare against a list of inappropriate objects
 - c. "Adult" and check if the Boolean is set to true (correct answer)**
 - d. "Inappropriate" and check if the Boolean is set to true
3. All objects listed in the image analysis response are guaranteed to be present in the image.
Select the correct answer.
- a. True, but they may be cut off
 - b. True, if they aren't in the image they won't be listed
 - c. False, objects cannot be detected in images using AI Vision
 - d. False, each image has a confidence score associated with it (correct answer)**
4. What are some examples of how Optical Character Recognition could be used in real world applications?
Select all that apply.
- a. Extract text from a handwritten note (correct answer)**
 - b. Transcribe text from one script to another (e.g. Japanese script to Latin script)
 - c. Extract text from product label (correct answer)**
 - d. Extract words from audio files
5. What API does AI Vision offer for reading text?
Select the correct option.
- a. The Vision API
 - b. The Image Analysis API (correct answer)**
 - c. The Language Reader API
 - d. The Text API
6. When using the Azure Image Analysis API, what do you specify when creating the ImageAnalyzer object?
Select all that apply.
- a. The endpoint and key for the Image Analysis service (correct answer)**
 - b. The features you want to retrieve from the image (correct answer)**
 - c. The image's file format
 - d. The color palette for the image

7. In addition to detecting a face in an image with a bounding box, what other capabilities does the Face service offer over the AI Vision service?

Select all that apply.

- a. **Facial recognition (correct answer)**
- b. **Comprehensive facial features (e.g. facial landmarks) (correct answer)**
- c. Age
- d. Gender

8. Using the AI Vision service, how can you get information (e.g. age and gender) for more than one person in a single image?

Select the correct answer.

- a. You must send the image through multiple times, blocking out all but one face for each submission
- b. **The AI Vision service will return a JSON object that contains an array of faces (correct answer)**
- c. The AI Vision service will return a JSON object that contains a list of face locations, and you must then specify which location you want additional information on
- d. You can only get information on one person per image

9. Can you use Face services to verify that two images are of the same person anonymously?

Select the correct answer.

a. Yes, the Face service assigns a unique ID to each face. This can be used to verify without needing additional identifiable information (correct answer)

b. Partially, while the Face service can anonymously detect that two faces are similar, it cannot verify that they are the same without a human verifying additional personally identifiable information

c. No, while the Face service is able to verify that two faces are the same, it can only do this by also storing additional identifiable information (e.g. name)

d. No, the Face service isn't able to verify that two faces are the same

10. Can Face services be used to identify a criminal based on a sketch made by a sketch artist?

Select the correct answer.

a. Yes, with a combination of Face services and Custom Vision a face-to-sketch comparison is possible

b. Yes, and this is commonly used

c. No, while it is possible to compare a face with a sketch, it is not permitted within the EULA

d. No, even if it were possible to have an accurate comparison, it is not permitted within the EULA (correct answer)

11. "Facial landmark location" is a feature of Face services that identifies the location in the image of what?

Select the correct answer.

- a. The landmarks that are near to the face (e.g. if a person is standing in front of the Eiffel Tower)
- b. Faces that are on popular landmarks (e.g. the Statue of Liberty)
- c. The specific facial features of a face (e.g. the tip of the nose) (correct answer)**
- d. The boundaries of the face (e.g. the top, sides, and bottom of each face in the image)

12. What is the correct sequence of steps for training a facial recognition model using Azure AI Services for persistent face recognition?

Select the correct answer.

- a. Add detected faces to a Person Group, create a Person for each individual, train the model, and then verify the identity of detected faces
- b. Create a Person Group, add a Person for each individual, add detected faces to each person, train the model, and then identify individuals in images (correct answer)**
- c. Train the model first, then add detected faces to a Person Group, create a Person for each individual, and finally identify individuals in images
- d. Create a Person Group, add detected faces to each person, train the model, and then use the model for both identifying and verifying individuals in new images

Custom vision

1. What are all the approaches to train a custom image classification model?

Select the correct answer.

- a. The Custom Vision Studio, the Custom Vision REST API or SDK, or a combination of both approaches. (correct answer)**
- b. Only the Vision Studio
- c. The Vision API or SDK, and the Azure Machine Learning portal
- d. The Azure portal or the Vision Studio

2. If a model has been trained to allow each image to be associated with multiple labels, what type of classification is that considered?

Select the correct answer.

- a. Multiclass classification
- b. Multi-image classification

c. Multilabel classification (correct answer)

d. Multi-object classification

3. How does Object Detection differ from Image Classification?

Select the correct answer.

a. Object Detection uses bounding boxes around objects within an image, while Image Classification will apply a label to the entire image. (correct answer)

b. Image Classification uses bounding boxes around objects within an image, while Object Detection will apply a label to the entire image

c. Object Detection can only identify objects based on an input model you provide while Image Classification uses complex AI to determine the objects

d. Image Classification can only identify objects based on an input model you provide while Object Detection uses complex AI to determine the objects

6. Which of the following sequences correctly represents the process for training a custom vision model to predict labels for new images?

Select the correct answer.

a. Label your images, upload your images, train your model, query your model with new images

b. Train your model, upload your images, label your images, query your model with new images

c. Upload your images, train your model, label your images, query your model with new images

d. Upload your images, label your images, train your model, query your model with new images (correct answer)**Open-ended questions**

1. What are some examples of information that can be extracted from images using AI Vision?

(Answer: Should include description and tag generation, object detection, face detection, image metadata/color/type analysis, category identification, brand detection, moderation rating, optical character recognition, smart thumbnail generation.)

2. A photo storage app wants to automatically feature a 100x100 pixel size image of the first image in every folder a user has. Can they do this using AI Vision? If so, how? If not, why not?

(Answer: Should include a brief explanation of the smart-cropped thumbnail feature that uses image analysis to determine a region of interest in an image, make it the focus, and crop to the appropriate size.)

3. In an attempt to store the handwritten notes of an English department at a University, you're tasked with creating an application that will allow researchers to take a picture of the documents and store the text digitally. Which API would you use and why?

(Answer: Should include the Read API because it can read handwritten text in English)

4. Describe one detail besides the actual text in the JSON response of the Image Analysis API.

(Answer: Should include one of the following: page, angle, bounding box, size, appearance style, confidence score)

5. Describe one consideration for the ethical use of face analysis in applications.

(Answer: Should include either Data privacy and security, Transparency, or Fairness and inclusiveness. For Data privacy and security, might mention that facial data is personally identifiable and should be considered sensitive and private. For Transparency, might mention that users need to be informed of how their facial data will be used and who will have access to it. For fairness and inclusiveness, might include that a face-based system cannot be user in a manner that is prejudicial to individuals based on their appearance, or to unfairly target individuals.)

6. A theme park takes 20 images of each rider while they ride down a roller coaster. Rather than display all 20 to the customers, they decide to use Face service to show the best 5. What functionality should they focus on for choosing the top 5 images??

(Answer: Should likely mention the Face Attribute Analysis, specifically head pose, blur, noise, and occlusion. Might also mention emotions and smile.)

7. A company wants to train a facial recognition model on their authorized employees for a particular lab. Describe the steps for creating this application using the Face services.

(Answer: Should mention creating a Person Group for the authorized employees, creating a Person within that Person Group for each individual you want represented, and then adding detected faces from multiple images of that individual to the Person. After this, they should train and deploy the model.)

8. A hospital wants to verify that the guest who checked in at reception is the same guest who entered a particular room. What should you do to solve this challenge?

(Answer: Should mention verifying the face detected at the room with the face ID of the face detected at reception.)

9. You need to create a facial recognition solution to identify specific employees at a company. Should you use Custom Vision or Face services for your solution? Why?

(Answer: Face because with Face services you can create a facial recognition solution.)

Custom vision

1. What is the difference between a model trained for multiclass classification versus multilabel classification?

(Answer: Models can be trained for multiclass classification (in other words, there are multiple classes, but each image can belong to only one class) or multilabel classification (in other words, an image might be associated with multiple labels).)

2. What is the class label for an image classification project?

(Answer: Image classification is an AI vision technique in which a model is trained to predict a class label for an image based on its contents. Usually, the class label relates to the main subject of the image.)

3. What is an example of Object Detection used at an in-person store?

(Answer: An AI-enabled checkout system in a grocery store might need to identify the type and location of items being purchased by a customer.)

4. What is the most significant difference between training an image classification model and training an object detection model?

(Answer: While image classification requires one or more tags that apply to the whole image, object detection requires that each label consists of a tag and a region that defines the bounding box for each object in an image. The Vision Studio provides a graphical interface that you can use to label your training images.)

5. A company wants to publish an app that allows customers to take a picture of individual pieces of fruit and the app will identify the type of fruit. Describe how you would use Custom Vision to do this.

(Answer: Should include a description of using the portal to upload images of fruit, tag the images, train and evaluate a classification model, test the model, and then publish the trained model that can be called from an API.)

5: Develop AI information extraction solutions

Multiple choice questions

Document intelligence

1. What is one reason why an intake desk might want to use the Azure AI Document Intelligence service instead of manually re-entering data from intake forms?

Select the correct answer.

a. Azure AI Document Intelligence service provide the building blocks for automation by using intelligent services to extract data at scale. (correct answer)

b. Azure AI Document Intelligence service will not make mistakes that are often made by humans when re-entering data manually

c. Azure AI Document Intelligence service can replace people at intake desks, reducing the cost of running a business

d. Azure AI Document Intelligence service should only be used on typed forms, and so wouldn't be useful at an intake desk

2. Azure AI Document Intelligence can be used on what formats?

Select all that apply.

a. JPG or PNG (correct answer)

b. PDF (correct answer)

c. TIFF (correct answer)

d. DOCX (correct answer)

3. What is an example of how the Business Cards prebuilt Azure AI Document Intelligence model could be used?

Select the correct answer.

a. Processing paperwork for each card owner in real time

b. Automatically creating contacts from images to integrate data with a CRM (correct answer)

c. Analyzing consumer behavior with business card drop offs

d. Improving business card design through comparisons of collected cards

4. What steps are involved in training a custom model for document intelligence in Azure?

Select the correct answer.

a. Connect to blob storage, auto label using an existing dataset, manually label all fields and files, and then train the model using the auto-generated Model ID

b. Create a project, upload training files or connect to blob storage, add data types, label data by selecting words and assigning fields, use layout and auto label to assist, and then train the model to obtain a Model ID (correct answer)

c. Upload training files to an existing project, label data using only the auto label feature, ignore manual selection of words, and train the model without generating a Model ID

d. Train the model first with a prebuilt model, then create a project, and finally label the data by selecting words in the document after training is complete

5. What information must be included in the API request when analyzing documents using a custom model in Azure Document Intelligence?

Select the correct answer.

a. Only the endpoint and key from the deployed resource

b. The endpoint, resource key, and the API version

c. The endpoint, key from the deployed resource, and the ID of the deployed custom model (correct answer)

d. A custom model ID only

Open-ended questions

1. What are the three services that Azure AI Document Intelligence provides that improve the service over a simple Read API?

(Answer: Should include Layout Service that returns information in bounding boxes, text content, tables, selection marks, and document structure, Prebuilt Models that currently support receipts, business cards, and invoices, and Custom Models that include unsupervised and supervised learning for training custom models for use in unlabeled and labeled forms.)

2. Describe an example of how the prebuilt Receipt model of the Azure AI Document Intelligence service could be used.

(Answer: Could include streamlining business expense reporting processes, automating auditing and accounting tasks, and analyzing consumer behavior and shopping trends.)

3. You are creating a receipt reader application for a customer and decide to use the prebuild Receipt model. What happens when you call the Analyze REST API function?

(Answer: Should include a description of how the result ID is returned and then used to call the Get Analyze Result function to actually retrieve the results. The results will be in one JSON response with

two nodes: readResults and documentResults, which contain all of the recognized text and the form-specific values.)