



Free Microsoft AB-731 Practice Questions

Shared by Garcia on 13-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page



Question 1

Question Type: MultipleChoice

Your company sells hiking and camping gear online. You need a generative AI solution that can interact with customers and ask questions about their needs. What should you include in the solution?

Options:

- A- a chatbot
- B- predictive AI
- C- computer vision
- D- a recommendation engine

Answer:

A

Explanation:

The requirement is an interactive generative AI experience that can converse with customers and ask clarifying questions (for example: "What climate are you hiking in?", "How many people will share the tent?", "What's your budget?", "Do you prioritize weight or comfort?"). The best solution component for that conversational, question-and-answer interaction is a chatbot (A), powered by a generative AI model.

A chatbot provides the dialog framework: maintaining conversational context across turns, prompting the user for missing requirements, and responding in natural language. This makes it suitable for customer support, guided shopping assistance, troubleshooting, and pre-sales Q&A--especially when customers don't know exactly what they need and benefit from a guided conversation.

The other options don't match the core requirement. Predictive AI (B) forecasts outcomes (like demand or churn) and isn't inherently conversational. Computer vision (C) analyzes images (like recognizing products from photos) and doesn't address asking questions in dialogue. A recommendation engine (D) can be useful in ecommerce, but it typically suggests items based on behavior or attributes; it doesn't by itself provide a conversational flow that asks users questions and adapts responses in natural language. In practice, you can combine a chatbot with a recommendation engine behind the scenes--but the "include in the solution" component that directly satisfies interactive questioning is the chatbot.

Question 2

Question Type: MultipleChoice

Your company is developing an AI-powered customer support agent. You need to ensure that the solution follows Microsoft responsible AI principles. Which two actions should you perform? Choose the two BEST answers. Each correct answer presents part of the solution.

Options:

- A- Retain all customer conversations.
- B- Test the agent to ensure that responses are inclusive and culturally sensitive.
- C- Enable the agent to operate independently.
- D- Ensure that the agent can be used for multiple purposes.
- E- Provide a clear disclaimer that users are interacting with an AI solution.

Answer:

B, E

Explanation:

To align an AI customer support agent with Microsoft's Responsible AI principles, two high-impact actions are fairness/inclusiveness validation and transparency to users. B is correct because testing for inclusive and culturally sensitive responses directly supports fairness and helps reduce harm. In practice, you evaluate responses across diverse user personas, languages/dialects, accessibility scenarios, and sensitive contexts. You look for biased assumptions, stereotyping, exclusionary language, and disparate quality of service. This also implies ongoing monitoring because model behavior can drift as prompts, knowledge sources, and user inputs evolve.

E is correct because a clear disclaimer supports transparency: customers should know they are interacting with an AI system, understand the type of assistance it can provide, and know what to do if the response is incorrect or they need a human. A disclosure is also a practical risk-control that reduces overreliance and sets expectations about limitations.

The other options are not best for Responsible AI alignment: A (retain all conversations) can conflict with privacy/data minimization; retention must be justified and governed, not automatic. C (operate independently) undermines accountability and human oversight. D (multiple purposes) increases scope and risk rather than improving responsible use.

Question 3

Question Type: MultipleChoice

You are exploring how Microsoft 365 Copilot uses Microsoft Graph to deliver AI-powered experiences. Which information in Microsoft Graph can Copilot use by default?

Options:

- A- emails, files, meetings, and chats in Microsoft 365
- B- data stored in a file share
- C- social media activity
- D- content from public websites

Answer:

A

Explanation:

Microsoft 365 Copilot is designed to work within the Microsoft 365 ecosystem and use organizational context that is already governed by Microsoft Entra ID, Microsoft 365 permissions, and compliance controls. By default, Copilot can use Microsoft Graph signals and content that exist in Microsoft 365 workloads the user already has access to---most commonly emails, files, meetings, and chats. That corresponds to A.

The key concept is permission-trimming: Copilot doesn't magically gain access to everything; it can only surface or use data that the signed-in user is permitted to access in Microsoft 365. This is what makes Copilot valuable for productivity scenarios---summarizing email threads, drafting replies, generating meeting recaps, creating documents from your files, or pulling context from Teams chats---because those artifacts are already part of daily work and already subject to tenant policies.

The other options are not "by default" Microsoft Graph content for Copilot: B (file shares) typically requires additional integration or migration into Microsoft 365 repositories or indexing via connectors; it's not inherently available. C is unrelated to Microsoft 365 tenant work data. D (public web) is not "information in Microsoft Graph"; web grounding is a separate capability and not the default Graph workload data source referenced here.

Question 4

Question Type: MultipleChoice

Your company plans to use generative AI to help project managers and engineers work with construction blueprints stored as PDF files. You need to recommend a generative AI solution that processes both images and text, summarizes building design, answers questions, and extracts information such as locations of electrical, heating, and plumbing systems. What should you recommend?

Options:

- A- a multi-modal solution
- B- an optical character recognition OCR solution
- C- a text completion solution
- D- a document summarization solution

Answer:

A

Explanation:

Construction blueprints in PDFs often contain a mix of text, symbols, linework, and diagrams. The requirements include understanding both visual layout (where systems are located) and textual annotations, producing summaries, and answering Q&A. That combination requires a multimodal generative AI approach---models that can reason over images and text together. Therefore, A is best.

OCR alone (B) can extract printed text, but it won't reliably interpret diagram geometry, symbols, or spatial relationships (e.g., "electrical riser is on the east core near gridline B-4"). Text completion (C) is too generic and doesn't address image understanding. Document summarization (D) is only one requirement (summary) and still depends on first extracting/understanding both visual and textual elements.

A multimodal solution can ingest the PDF pages as images (or rendered page images) plus extracted text, then answer questions grounded in both modalities. In practice, you may combine OCR and layout extraction with a multimodal LLM so the model can reference drawing regions, legends, callouts, and system diagrams to produce accurate explanations and field extractions.

Question 5

Question Type: MultipleChoice

You need to create a custom Azure Machine Learning model. The data used to train the model is consistent and uniform. What should you do first?

Options:

- A- Prepare the training data.
- B- Evaluate the model.
- C- Train the model.
- D- Tune hyperparameters.
- E- Deploy the model.

Answer:

A

Explanation:

Even when training data is already consistent and uniform, the first step in building a custom Azure Machine Learning model is still to prepare the training data. "Consistent" data reduces the amount of cleaning you may need, but preparation is broader than cleaning: you still must confirm the schema, validate data types, handle missing values (if any), ensure label quality (for supervised learning), select/engineer features, and split data into training/validation/test sets. Those actions determine whether training will be stable and whether evaluation metrics will be meaningful.

If you skip preparation and go directly to training (C), the model might learn from the wrong columns, inconsistent labels, or poorly partitioned data, producing misleading results. Evaluation (B) comes after training because you need a trained model to score and measure. Hyperparameter tuning (D) is an optimization activity that presupposes you already have a working training pipeline and a baseline model to improve. Deployment (E) is last, after you have validated performance and selected the model candidate.

Azure Machine Learning commonly operationalizes these steps through pipelines, where data preparation is a foundational stage that precedes training and evaluation (and can also be iterated as you refine features and quality).

Question 6

Question Type: MultipleChoice

Your company sells hiking and camping gear online. You need a generative AI solution that can interact with customers and ask questions about their needs. What should you include in the solution?

Options:

- A- a chatbot
- B- predictive AI
- C- computer vision
- D- a recommendation engine

Answer:

A

Explanation:

The requirement is an interactive generative AI experience that can converse with customers and ask clarifying questions (for example: "What climate are you hiking in?", "How many people will share the tent?", "What's your budget?", "Do you prioritize weight or comfort?"). The best solution component for that conversational, question-and-answer interaction is a chatbot (A), powered by a generative AI model.

A chatbot provides the dialog framework: maintaining conversational context across turns, prompting the user for missing requirements, and responding in natural language. This makes it suitable for customer support, guided shopping assistance, troubleshooting, and pre-sales Q&A--especially when customers don't know exactly what they need and benefit from a guided conversation.

The other options don't match the core requirement. Predictive AI (B) forecasts outcomes (like demand or churn) and isn't inherently conversational. Computer vision (C) analyzes images (like recognizing products from photos) and doesn't address asking questions in dialogue. A recommendation engine (D) can be useful in ecommerce, but it typically suggests items based on behavior or attributes; it doesn't by itself provide a conversational flow that asks users questions and adapts responses in natural language. In practice, you can combine a chatbot with a recommendation engine behind the scenes--but the "include in the solution" component that directly satisfies interactive questioning is the chatbot.

Question 7

Question Type: MultipleChoice

You have a business unit that uses an AI solution to process loan applications. You discover that the solution rejects the application of all applicants that are older than 60 years of age. Which Microsoft responsible AI principle is this violating?

Options:

- A- accountability
- B- transparency
- C- fairness
- D- reliability and safety

Answer:

C

Explanation:

This scenario is a clear violation of the fairness principle. Fairness in Microsoft's Responsible AI framework is about ensuring AI systems do not create unjustified bias or discriminatory outcomes---especially when decisions affect people's access to opportunities such as credit, employment, housing, or education. A rule or learned behavior that rejects all applicants over a certain age creates a systematic, categorical disadvantage for a protected demographic group and indicates a discriminatory decision boundary rather than an individualized assessment of creditworthiness.

Even if the model designers believed age correlates with risk, using a hard cutoff that rejects every applicant older than 60 is not an equitable approach. It suggests the model is either using age directly as a dominant feature or reflects biased training data/labels that encoded discriminatory outcomes. Fairness requires you to evaluate model outcomes across groups (for example, age brackets), measure disparate impact, and apply mitigations such as feature review (removing or constraining sensitive attributes), rebalancing training data, adjusting thresholds, or using fairness-aware training/evaluation methods. It also requires governance and review of high-stakes automated decisions.

The other principles are not the best match: transparency concerns explainability and user understanding, accountability concerns human oversight and ownership, and reliability and safety concerns consistent and safe operation. The core issue here is discriminatory treatment across an age group---fairness.

Question 8

Question Type: MultipleChoice

Your company uses a generative AI solution. You need to improve the quality of responses by using grounding. Which statement accurately describes how grounding improves accuracy and relevancy?

Options:

- A- references a diverse set of people, disciplines, and perspectives
- B- explains how and why AI models generate content
- C- anchors the responses in specific data sources
- D- specifies the strengths and weaknesses of the AI model

Answer:

C

Explanation:

Grounding is an AI solution pattern used to improve response quality by ensuring the model's output is based on trusted, relevant information provided at inference time, rather than relying only on what the model "remembers" from training. Therefore, C is correct: grounding anchors responses in specific data sources.

In practical deployments, grounding commonly uses retrieval (often called Retrieval Augmented Generation, or RAG) where the system first finds relevant content from approved sources---such as internal policy documents, product documentation, knowledge bases, or databases---and then includes that content in the prompt context sent to the model. Microsoft's guidance describes grounding data as information supplied at inference time to help responses become more accurate and relevant because the model is guided by authoritative, up-to-date content that may not have been part of original training.

The other options do not define grounding. A relates to inclusion practices and diversity considerations, which are important for responsible AI but are not what grounding means. B describes transparency/explainability concepts. D relates to model evaluation/communication of limitations. Grounding is specifically about tying outputs to known sources, which reduces hallucinations and improves business trust in the generated responses.

Question 9

Question Type: MultipleChoice

Your company is preparing to adopt Microsoft 365 Copilot and wants to follow Microsoft responsible AI principles. As a business leader, you propose establishing an AI governance council to ensure alignment with the responsible AI principles. What is the main purpose of the council? More than one answer choice may achieve the goal. Choose the BEST answer.

Options:

- A- to monitor user behavior and enforce compliance with internal IT policies
- B- to train employees on how to use Copilot features effectively
- C- to oversee implementation, manage technical performance, and ensure successful AI deployment
- D- to guide strategy, provide oversight, and ensure cross-functional alignment for responsible AI adoption

Answer:

D

Explanation:

An AI governance council (often called an "AI Council") exists primarily to set direction and provide cross-functional oversight so AI adoption stays aligned to the organization's values, risk posture, and Responsible AI commitments. That maps most directly to D. Microsoft's guidance on creating an AI Council describes leadership responsibilities such as defining and communicating the organization's AI vision, values, and policies, reviewing and approving AI use cases/projects, and coordinating with enablement and technical readiness teams to understand risks, issues, and opportunities. It also emphasizes representation across distinct functions (for example: senior leadership, legal, compliance, risk, ethics, data, technology, business, HR) to ensure governance decisions reflect a broad, accountable perspective.

The other options describe activities that may be supporting outcomes of governance, but they are not the council's primary purpose. A is narrow (IT policy enforcement/user monitoring) and is typically handled by security/compliance operations rather than the top-level governance body. B is user enablement/training (commonly owned by adoption/change management teams). C focuses on technical delivery and performance management (often owned by engineering/MLOps/service owners). The governance council's central value is strategic guidance + oversight + cross-functional alignment to ensure Responsible AI adoption is consistent, accountable, and sustainable across the business.

Question 10

Question Type: MultipleChoice

Your company uses a non-reasoning generative AI model to create textual content. You discover that the model's responses are inconsistent and do NOT meet expectations. You need to improve the prompts. What should you do? More than one answer choice may achieve the goal. Select the BEST answer.

Options:

- A- Provide the prompts with extensive examples of the expected output.
- B- Add the context, sources, and expectations to the prompts.
- C- Use technical terms in the prompts to enhance AI comprehension.
- D- Add only a single concise requirement to the prompts.

Answer:

A, B

Explanation:

When a non-reasoning generative AI model produces inconsistent outputs, the most reliable improvement is to make the prompt more specific, constrained, and demonstrative of what "good" looks like.

A is correct because adding high-quality examples is a form of few-shot prompting. Examples act like "training wheels" at inference time: they show the model the desired structure, tone, level of detail, formatting rules, and boundaries. This reduces ambiguity and variance, especially for tasks like marketing copy, summaries, policy text, or customer replies. The more your examples resemble real target outputs (including edge cases), the more consistent the model's completions become.

B is correct because adding context, relevant source material, and explicit expectations narrows the model's degrees of freedom. Including the intended audience, purpose, constraints (length, voice, banned claims), and trusted reference content (approved facts, product specs, policy excerpts) helps the model stay aligned and reduces hallucinations and off-brand language. This is also where you specify acceptance criteria such as "must include 3 bullet points," "use UK English," or "cite only provided text."

C is not best: technical jargon can confuse or bias output if it's not aligned to the task; clarity beats jargon. D is not best: a single concise requirement is usually under-specified and often increases variability.

Question 11

Question Type: MultipleChoice

Your company purchases Microsoft 365 Copilot for its sales department. The sales department needs to find and summarize information across internal documents quickly. From which two data sources can the sales department obtain results by default? (Select TWO.)

Options:

- A- an on-premises file share
- B- a custom customer relationship management (CRM) system
- C- Microsoft SharePoint
- D- Microsoft OneDrive
- E- Microsoft Sway

Answer:

C, D

Explanation:

By default, Microsoft 365 Copilot is grounded in your organization's Microsoft 365 data through Microsoft Graph, and it respects the user's existing permissions. For "find and summarize information across internal documents," the most direct default document repositories in Microsoft 365 are SharePoint (team sites and shared libraries) and OneDrive (a user's work files). That is why C (Microsoft SharePoint) and D (Microsoft OneDrive) are the correct selections. Microsoft explicitly describes Copilot as accessing organizational content via Microsoft Graph, including user documents and related work content.

The other options are not "by default" sources. A (on-premises file share) is not automatically part of Microsoft Graph unless you integrate/migrate content or use connectors to make it discoverable in Microsoft 365 experiences. B (custom CRM) similarly requires an integration approach (for example, Microsoft 365 Copilot connectors / Graph connectors) to index and expose that data for Copilot to use. E (Microsoft Sway) is not a primary default content source for Copilot's document grounding and is not typically referenced as a core internal document repository compared to SharePoint/OneDrive.

Question 12

Question Type: MultipleChoice

Your company discovers that several employees use personal ChatGPT accounts to assist with work tasks. You are concerned about proprietary data being shared externally. You need to evaluate the business value of rolling out Microsoft 365 Copilot. Which capability is a key benefit of using Copilot instead of a personal ChatGPT account?

Options:

- A- analyzing and producing reports based on complex data
- B- generating ideas and solving issues
- C- drafting documents, emails, presentations, and marketing materials
- D- accessing internal data in accordance with existing Microsoft 365 policies

Answer:

D

Explanation:

The core business concern in the scenario is data leakage---employees using consumer tools where corporate data could be pasted, stored, or processed outside the organization's governance boundary. The key differentiator of Microsoft 365 Copilot is that it's designed to work inside your Microsoft 365 tenant and to respect the organization's existing security, compliance, identity, and data access controls. Therefore, D is the best answer: Copilot accesses internal work data (Microsoft Graph-connected content such as mail, files, chats, meetings) in accordance with existing Microsoft 365 policies and permissions---meaning it can only surface content the user is already allowed to access, and it operates under enterprise-grade controls (authentication, auditing, compliance boundaries, and admin governance).

Options B and C describe general generative AI capabilities that personal ChatGPT can also provide (brainstorming, drafting, rewriting). A can be done in multiple tools as well, and it is not the primary "enterprise value" difference tied to the stated risk. The scenario's driver is governance: reducing the likelihood of proprietary data leaving controlled systems while still enabling productivity. Rolling out Copilot addresses that by providing "work-safe" AI anchored to organizational content and managed through the same tenant controls your company already uses.



To Get Premium Files for AB-731 Visit

<https://www.p2pexams.com/products/ab-731>

For More Free Questions Visit

<https://www.p2pexams.com/microsoft/pdf/ab-731>

20%
DISCOUNT

P2P
exams