

Microsoft.AB-731.v2026-09-03.q37

Exam Code:	AB-731
Exam Name:	AI Transformation Leader
Certification Provider:	Microsoft
Free Question Number:	37
Version:	v2026-09-03
# of views:	144
# of Questions views:	370
https://www.freepdfdumps.com/Microsoft.AB-731.v2026-09-03.q37.html	

NEW QUESTION: 1

- For each of the following statements, select Yes if the statement is true. Otherwise, select No . NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Prompt engineering changes how a generative AI model was trained.	☐	☐
Prompt engineering focuses on designing clear, concise, and context-rich instructions.	☐	☐
Effective prompt engineering involves maximizing the number of tokens used in each request.	☐	☐

Answer:

Answer Area

Statements	Yes	No
Prompt engineering changes how a generative AI model was trained.	☐	☒
Prompt engineering focuses on designing clear, concise, and context-rich instructions.	☒	☐
Effective prompt engineering involves maximizing the number of tokens used in each request.	☐	☒

Explanation:

Answer Area

- * Prompt engineering changes how a generative AI model was trained. Answer: No
- * Prompt engineering focuses on designing clear, concise, and context-rich instructions. Answer: Yes
- * Effective prompt engineering involves maximizing the number of tokens used in each request. Answer: No
- * No - Prompt engineering does not modify model weights or retrain the model. It is an inference-time technique : you steer outputs by improving the instructions and context you send to the model. Changing how a model was trained would involve pretraining, fine-tuning, or other training methods- separate from prompt engineering.
- * Yes - The primary goal of prompt engineering is to reduce ambiguity and variability by providing clear instructions , the right context , and explicit output constraints . This often includes specifying role and task, providing necessary facts or grounding text, defining format (bullets, JSON, headings), and adding examples (few-shot) when helpful. Well-constructed prompts improve consistency, relevance, and usefulness of outputs.
- * No - Good prompt engineering does not mean "use as many tokens as possible." In fact, using unnecessary tokens can increase cost and may degrade quality by adding noise. Effective prompts are typically as short as possible but as long as necessary : enough context to achieve accuracy and alignment, but not bloated. Token discipline matters because most pricing is token-based and long contexts can dilute attention over what's most important.

NEW QUESTION: 2

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

Answer Area		
Statements	Yes	No
Microsoft 365 Copilot can help users compose professional emails and summarize conversations.	☐	☐
When prompted, Microsoft 365 Copilot can autonomously update a spreadsheet that contains sales data.	☐	☐
Microsoft 365 Copilot can summarize recent updates in business documents, spreadsheets, and presentations.	☐	☐

Answer:

Answer Area

Statements

Microsoft 365 Copilot can help users compose professional emails and summarize conversations.

Yes

☒

No

☐

When prompted, Microsoft 365 Copilot can autonomously update a spreadsheet that contains sales data.

☐☒

Microsoft 365 Copilot can summarize recent updates in business documents, spreadsheets, and presentations.

☒☐

Microsoft

Explanation:

Answer Area

Statements

Microsoft 365 Copilot can help users compose professional emails and summarize conversations.

Yes

☒

No

☐

When prompted, Microsoft 365 Copilot can autonomously update a spreadsheet that contains sales data.

☐☒

Microsoft 365 Copilot can summarize recent updates in business documents, spreadsheets, and presentations.

☒☐

Microsoft

The correct selections are Yes, No, Yes. Microsoft 365 Copilot helps users create, rewrite, summarize, and analyze content in Microsoft 365 apps such as Outlook, Word, Excel, PowerPoint, and Teams. Therefore, it can help compose professional emails and summarize conversations. It can also summarize recent updates across business content such as documents, spreadsheets, and presentations when the user has access to that content. The second statement is false because Copilot is not generally described as autonomously updating business spreadsheets without user control. It can assist with analysis, formulas, summaries, and suggested changes, but autonomous modification of sales data would require governed user action, permissions, and potentially workflow controls. Copilot is an assistant, not an unrestricted autonomous spreadsheet operator.

NEW QUESTION: 3

Select the answer that correctly completes the sentence.

When Microsoft 365 Copilot accesses company data using the principle of least privilege, this is an example of a control.

data classification
data compliance
data loss prevention (DLP)
data security



Answer:

When Microsoft 365 Copilot accesses company data using the principle of least privilege, this is an example of a control.

data classification
data compliance
data loss prevention (DLP)
data security



Explanation:

The correct answer is data security. The principle of least privilege means users, apps, and services receive only the minimum access required to perform their work. Microsoft 365 Copilot respects existing Microsoft 365 permissions and does not grant users new access to content they are not already authorized to view. That behavior is a security control because it helps prevent unauthorized access to company information. Data classification is about labeling or categorizing information by sensitivity. Data compliance focuses on meeting policy, legal, or regulatory requirements. Data loss prevention helps prevent sensitive information from being shared inappropriately. Least privilege is primarily an access-control practice, so the best category is data security.

NEW QUESTION: 4

- Select the answer that correctly completes the sentence.

An organization that runs continuous, large-scale workloads with Azure OpenAI models should choose the _____ pricing model.

Answer Area

An organization that runs continuous, large-scale workloads with Azure OpenAI models should choose the pricing model.

Batch API
Provisioned (PTUs)
Standard (On-Demand)



Answer:



Explanation:

Provisioned (PTUs)

For continuous, large-scale workloads, the key needs are predictable throughput, consistent latency, and cost/performance stability . Azure OpenAI Provisioned Throughput Units (PTUs) are designed for this scenario because you reserve model capacity to meet sustained demand. This reduces the risk of variability that can occur with purely on-demand usage during peak periods and provides a more predictable operating model when the workload is always "on" and high volume.

Standard (On-Demand) is best when usage is variable or you're starting small (PoCs, pilots, spiky workloads), because it is flexible pay-as-you-go but can be less predictable at sustained scale. Batch API is optimized for asynchronous, non-interactive processing where you can tolerate delayed results (for example, large offline summarization jobs), not for always-on, real-time interactions at scale.

Therefore, for continuous high-volume production workloads, Provisioned (PTUs) is the best pricing model choice.

NEW QUESTION: 5

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?

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

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 6

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?

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

Answer: A ([LEAVE A REPLY](#))

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.

NEW QUESTION: 7

Select the answer that correctly completes the sentence.

Answer Area

An organization's AI champions program includes

- users who advocate for the value of AI by learning, sharing, and identifying use cases.
- technical teams responsible for designing advanced analytics solutions to track adoption.
- users who advocate for the value of AI by learning, sharing, and identifying use cases.
- users who are dedicated to building the infrastructure required to support AI across the organization.
- users who come together and define the organization's overall AI strategy.**

Answer:

Answer Area

An organization's AI champions program includes

- users who advocate for the value of AI by learning, sharing, and identifying use cases.
- technical teams responsible for designing advanced analytics solutions to track adoption.
- users who advocate for the value of AI by learning, sharing, and identifying use cases.
- users who are dedicated to building the infrastructure required to support AI across the organization.
- users who come together and define the organization's overall AI strategy.**

Explanation:

Answer Area

An organization's AI champions program includes

- users who advocate for the value of AI by learning, sharing, and identifying use cases.

An AI champions program includes users who advocate for AI by learning, sharing, and identifying use cases. Champions are not primarily the infrastructure team, analytics team, or executive strategy group. Their value is adoption enablement. They help colleagues understand AI tools, share practical usage patterns, collect feedback, and discover department-specific opportunities where AI can improve productivity or decision-making. Strategy and governance are usually handled by leadership, AI councils, legal, risk, compliance, security, and technical stakeholders. Infrastructure is handled by IT and engineering teams. Champions sit closer to end users and help translate the AI strategy into practical daily usage. This makes them critical for scaling proficiency and adoption across the organization.

NEW QUESTION: 8

Your company creates a custom Azure Machine Learning model that uses a generative AI assistant. The model initially delivers strong results. However, six months later, the model predictions become noticeably less accurate. What is a possible cause of the issue?

- A. The input data changed over time.
- B. The model requires additional compute resources.
- C. The model was trained incorrectly.

Answer: A ([LEAVE A REPLY](#))

A common reason models degrade after being successful in production is data drift (also called concept drift). Over time, the distribution of input data changes -for example, customer behavior shifts, product catalog changes, seasonality changes, new categories appear, sensors get recalibrated, or business processes evolve. When the model sees data that differs from what it was trained on, its predictions can become less accurate. This is exactly what option A describes and is the most likely "six months later" cause.

Option B is not a primary explanation for reduced predictive accuracy. More compute can improve throughput /latency, but it does not inherently improve correctness of predictions. If anything, compute constraints typically cause timeouts or slower responses, not a systematic accuracy drop.

Option C (trained incorrectly) would usually manifest earlier-poor performance from the start-unless the "incorrectness" is that the model was trained on a snapshot that later became stale (which again maps back to drift). The correct operational response is to monitor for drift, validate performance regularly, and retrain /refresh the model using newer representative data and updated features/labels.

NEW QUESTION: 9

Which statement accurately describes the difference between a pretrained generative AI model and a fine- tuned generative AI model?

- A. A pretrained model requires labeled data, while a fine-tuned model does not.
- B. A pretrained model is faster to train than a fine-tuned model because the pretrained model uses fewer parameters.
- C. A pretrained model is trained on broad datasets, while a fine-tuned model is adapted to perform well on a narrower, domain-specific dataset.
- D. A pretrained model is optimized for a specific task, while a fine-tuned model is designed for general- purpose use.

Answer: C (LEAVE A REPLY)

A pretrained generative AI model is trained initially on a large, broad, and diverse dataset so it learns general language (or multimodal) patterns and capabilities. Fine-tuning then takes that pretrained base and performs additional training on a smaller, task- or domain-specific dataset to specialize behavior- improving performance for a particular use case, tone, style, or domain knowledge representation. That is exactly what option C states, making it the correct answer.


Option A is incorrect because both pretraining and fine-tuning may use labeled or unlabeled data depending on the technique; the distinction is not "labeled vs. unlabeled." Option B is incorrect because a pretrained model is not "faster to train" due to fewer parameters; pretraining is typically the most compute-intensive phase precisely because it's done at large scale, while fine-tuning is smaller but still trains the same model architecture. Option D is reversed: the pretrained model is the general-purpose foundation, while the fine- tuned model is the specialized variant for a specific task or dataset.

NEW QUESTION: 10

- Select the answer that correctly completes the sentence.
The Analyst agent in Microsoft 365 Copilot _____.



Answer:



Explanation:
uses structured data and provides insights by using text, charts, tables, and other visuals.

The Analyst agent in Microsoft 365 Copilot is positioned as a "data analysis" reasoning agent that helps users work through structured information (for example, tables, spreadsheets, and other dataset-like inputs) and then produces analytical outputs . The best completion is the option stating it "uses structured data and provides insights by using text, charts, tables, and other visuals," because that describes the hallmark outcome of analyst-style work: summarizing patterns, highlighting key drivers, and presenting results in formats that business users can act on. Analyst-style assistance typically includes exploring the data, identifying trends and anomalies, comparing segments, and explaining findings clearly-often accompanied by tables and visual representations that make the insights easier to consume. The other dropdown options align to different use cases: "compiles background research for a new market or initiative" describes a research-oriented agent, "generate audio summaries" is a media summarization function, and "answer employee FAQs" describes a conversational knowledge assistant. Analyst is the one most directly associated with structured-data interpretation and producing a mix of narrative plus analytical artifacts (tables/charts) to communicate conclusions.

NEW QUESTION: 11

Your company plans to adopt AI across multiple business units. You need to ensure that all AI projects align with the company's business strategy and are implemented responsibly. What is the best approach to achieve the goal? More than one answer choice may achieve the goal. Select the BEST answer.

A. Allow each department to deploy its own AI tools and workflows.

B. Delegate AI decision-making to the company's IT department.

C. Outsource AI development to an external vendor.

D. Establish an AI council to provide guidance, oversight, and coordination.

Answer: D (LEAVE A REPLY)

When AI adoption spans multiple business units, the primary risk is fragmented delivery: inconsistent standards, duplicated spend, uneven risk controls, and misalignment with enterprise strategy. Establishing an AI council (D) is the best approach because it creates a cross-functional governance mechanism that aligns AI initiatives to business priorities while enforcing Responsible AI practices consistently. An AI council typically includes senior stakeholders from business leadership, IT, security, legal, compliance, privacy, data, and HR. Its role is to define AI principles and guardrails, approve high-impact use cases, set policy for data usage and access, establish evaluation and monitoring requirements, and coordinate change management and training. This also enables portfolio management-deciding which projects to prioritize, reuse, or stop-so AI investments map to measurable business outcomes. The other options are weaker: A encourages siloed deployments and inconsistent risk management. B centralizes too narrowly in IT; Responsible AI requires broader accountability than a single function. C can help delivery capacity but does not replace internal governance; vendors still need direction, controls, and oversight from the organization.

NEW QUESTION: 12

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Allowing AI models to make autonomous decisions supports the Microsoft responsible AI principle of accountability.	☐	☐
Regularly testing AI models for fairness and inclusiveness helps ensure they align with Microsoft's Responsible AI principles.	☐	☐
Protecting user data and limiting access to personal information supports the Microsoft responsible AI principles of privacy and security.	☐	☐

Answer:

Answer Area

Statements	Yes	No
Allowing AI models to make autonomous decisions supports the Microsoft responsible AI principle of accountability.	☐	☒
Regularly testing AI models for fairness and inclusiveness helps ensure they align with Microsoft's Responsible AI principles.	☒	☐
Protecting user data and limiting access to personal information supports the Microsoft responsible AI principles of privacy and security.	☒	☐

Explanation:

Answer Area

* Allowing AI models to make autonomous decisions supports the Microsoft responsible AI principle of accountability. Answer: No

* Regularly testing AI models for fairness and inclusiveness helps ensure they align with Microsoft's Responsible AI principles. Answer: Yes

* Protecting user data and limiting access to personal information supports the Microsoft responsible AI principles of privacy and security. Answer: Yes Microsoft's Responsible AI principles emphasize that people and organizations must remain accountable for AI systems and their outcomes. Accountability is strengthened by governance, human oversight, clear ownership, auditability, and processes to review and address issues- not by letting models make unchecked autonomous decisions. Therefore, statement 1 is No : increasing autonomy can actually increase risk unless paired with human-in-the-loop controls and clear escalation paths, because accountability requires clear responsibility for decisions and impacts.

Statement 2 is Yes because fairness and inclusiveness are explicitly supported through ongoing evaluation.

Regular testing helps detect disparate impact, performance gaps across user groups, and unintended bias introduced by data drift or changes in usage patterns. It's not a one-time activity; it's continuous assurance that the system behaves appropriately as conditions change.

Statement 3 is Yes because privacy and security are directly supported by protecting personal/sensitive data, enforcing least privilege access, and implementing controls such as data loss prevention, encryption, access logging, and strong identity governance. Limiting access to personal information reduces exposure and supports compliance obligations while aligning with privacy-by-design and secure-by-design expectations for AI-enabled solutions.

NEW QUESTION: 13

Your company plans to implement a proof of concept PoC agent that uses Azure OpenAI. The solution must start small and provide flexibility to scale usage as demand grows. Which pricing model should you use?

A. Microsoft 365 Copilot

B. Batch API

C. Provisioned PTUs

D. Standard On-Demand

Answer: (SHOW ANSWER)

For a proof of concept , the key requirements are low commitment , quick start , and the ability to scale up or down as you learn what real usage looks like. Azure OpenAI Standard On-Demand pricing is designed for exactly that: you pay per token consumed (input and output) on a pay-as-you-go basis, which makes it ideal when demand is uncertain or variable-typical in early pilots and PoCs.

By contrast, Provisioned (PTUs) is best when you have well-defined, predictable throughput and latency requirements -usually a more mature, production workload. PTUs involve reserving model processing capacity to achieve consistent performance and more predictable costs, which is usually premature for a PoC where actual traffic patterns are not yet known.

Batch API is optimized for asynchronous high-volume jobs with a target turnaround (for example, up to 24 hours) and discounted pricing. That's great for offline processing, but it does not match an interactive "agent" PoC that typically needs near-real-time responses and iterative testing.

Microsoft 365 Copilot is a separate SaaS licensing model and is not the Azure OpenAI pricing model for building your own agent solution.

NEW QUESTION: 14

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?

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 (LEAVE A REPLY)

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.

NEW QUESTION: 15

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.

Answer Area

Statements	Yes	No
You can use Azure Language in Foundry Tools to analyze the sentiment of customer reviews.	☐	☐
You can use Azure Language in Foundry Tools to translate internal reports into multiple languages.	☐	☐
You can use Azure Language in Foundry Tools to extract text from scanned documents.	☐	☐

Answer:

Answer Area

Statements	Yes	No
You can use Azure Language in Foundry Tools to analyze the sentiment of customer reviews.	☒	☐
You can use Azure Language in Foundry Tools to translate internal reports into multiple languages.	☐	☒
You can use Azure Language in Foundry Tools to extract text from scanned documents.	☐	☒

Explanation:
The answer Area

* You can use Azure Language in Foundry Tools to analyze the sentiment of customer reviews. Answer: Yes

* You can use Azure Language in Foundry Tools to translate internal reports into multiple languages. The answer: No

* You can use Azure Language in Foundry Tools to extract text from scanned documents. Answer: No Azure Language is designed for natural language processing (NLP) over text that is already machine- readable. That includes capabilities like sentiment analysis , key phrase extraction, entity recognition, summarization, and classification. Therefore, statement 1 is Yes : sentiment analysis of customer reviews is a standard NLP workload where the service scores text as positive/negative/neutral (and often provides confidence scores), helping organizations quantify customer satisfaction and detect recurring issues. Statement 2 is No because translation is typically handled by a dedicated translation capability (commonly delivered as a separate translator service) rather than the core "Language" NLP features. While translation is an AI language workload, it's not what the Azure Language service is primarily used for in this context; the expected Microsoft service choice for multi-language translation is the translator capability, not Azure Language. Statement 3 is No because extracting text from scanned documents is OCR (optical character recognition), which is a computer vision/document processing function. OCR is delivered through Azure Vision and/or Azure Document Intelligence , which can read printed/handwritten text from images and PDFs and return structured output. Azure Language can analyze extracted text after OCR, but it does not perform the image-to-text extraction step itself.

NEW QUESTION: 16

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
A generative AI solution is well-suited to predict next-quarter sales trends.	☐	☐
A generative AI solution can summarize lengthy policy documents.	☐	☐
A generative AI solution can create product descriptions from product specifications.	☐	☐

Answer:

Statements	Yes	No
A generative AI solution is well-suited to predict next-quarter sales trends.	☒	☐
A generative AI solution can summarize lengthy policy documents.	☒	☐
A generative AI solution can create product descriptions from product specifications.	☒	☐

Explanation:
Answer Area

* A generative AI solution is well-suited to predict next-quarter sales trends. Answer: No

* A generative AI solution can summarize lengthy policy documents. Answer: Yes

* A generative AI solution can create product descriptions from product specifications. Answer: Yes

- * No - Predicting next-quarter sales trends is primarily a forecasting/predictive analytics problem. Microsoft differentiates predictive AI (forecasting outcomes from historical patterns) from generative AI (creating content like text, images, or code). While you can use LLMs to assist analysts (explain trends, draft narratives), the core forecasting model is typically traditional ML/time-series methods rather than generative AI as the main engine.
- * Yes - Summarization is a classic, high-value generative AI capability. Given a long policy, an LLM can compress it into executive summaries, key obligations, risks, and action items, often with formatting constraints (bullets, sections, "do/don't" lists). Microsoft highlights summarization and analysis as common generative AI use cases in business contexts.
- * Yes - Generative AI is well-suited to transform structured inputs (features/specs) into natural- language outputs (product descriptions). This is straightforward "content generation," where you control tone, length, and required fields (benefits, differentiators, disclaimers). Microsoft also points to generating product descriptions and similar marketing/customer-facing text as a practical generative AI scenario.

Valid AB-731 Dumps shared by Actual4test.com for Helping Passing AB-731 Exam! Actual4test.com now offer the **newest AB-731 exam dumps**, the Actual4test.com AB-731 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AB-731 dumps with Test Engine here: https://www.actual4test.com/AB-731_examcollection.html (117 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 17

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Microsoft Foundry provides a unified platform for developers and data professionals to create custom end-to-end AI solutions.	☐	☐
Microsoft Foundry provides a unified platform for low-code developers and business users to create custom end-to-end AI solutions.	☐	☐
You need a Microsoft 365 Copilot license to access Microsoft Foundry services.	☐	☐

Answer:

Answer Area

Statements	Yes	No
Microsoft Foundry provides a unified platform for developers and data professionals to create custom end-to-end AI solutions.	☒	☐
Microsoft Foundry provides a unified platform for low-code developers and business users to create custom end-to-end AI solutions.	☐	☒
You need a Microsoft 365 Copilot license to access Microsoft Foundry services.	☐	☒

Explanation:

Answer Area

- * Microsoft Foundry provides a unified platform for developers and data professionals to create custom end-to-end AI solutions. Answer: Yes
- * Microsoft Foundry provides a unified platform for low-code developers and business users to create custom end-to-end AI solutions. Answer: No
- * You need a Microsoft 365 Copilot license to access Microsoft Foundry services. Answer: No
- * Yes - Microsoft Foundry is positioned by Microsoft as a unified AI platform for building, deploying, and scaling AI apps and agents, with enterprise governance in a single environment. Microsoft explicitly describes Foundry as enabling developers (and, in broader descriptions, data professionals /enterprises) to build and operate AI solutions end-to-end, which aligns with the statement.

- * No - While Foundry can streamline AI development workflows and may include experiences that reduce complexity, the statement specifically targets low-code developers and business users creating custom end-to-end AI solutions. In Microsoft's ecosystem, that "low-code/business-user" agent/app- building lane is more directly associated with tools like Copilot Studio/Power Platform , whereas Foundry is primarily presented as the platform for developers and technical teams building and governing AI solutions. Therefore, as written, the statement is not the best characterization.
- * No - Access to Microsoft Foundry services is not gated by a Microsoft 365 Copilot license. Foundry is part of the Azure AI platform experience (formerly Azure AI Studio) and is accessed through Azure, using Azure subscription-based service consumption. Microsoft 365 Copilot licensing is for Microsoft 365 Copilot experiences, not a prerequisite to use Foundry services.

NEW QUESTION: 18

Select the answer that correctly completes the sentence.

When a generative AI model produces output that seems realistic but contains incorrect information, the behavior is known as _____.

Answer Area

When a generative AI model produces output that seems realistic but contains incorrect information, the behavior is known as

data leakage.

model inaccuracy.

overreliance.

prompt injection.

Answer:

Answer Area

When a generative AI model produces output that seems realistic but contains incorrect information, the behavior is known as

data leakage.

model inaccuracy. I

overreliance.

prompt injection.

Explanation:

model inaccuracy

The scenario describes a model producing plausible-sounding content that is factually wrong -a common generative AI failure mode often referred to as a "hallucination." Since "hallucination" is not offered in the dropdown, the best matching choice is model inaccuracy because the core problem is that the model's output is incorrect even though it appears confident and coherent.

The other options do not fit the definition of the behavior: data leakage is about sensitive information being exposed (for example, proprietary prompts, secrets, or personal data). Prompt injection is an attack technique where a user tries to override system instructions or cause unsafe actions. Overreliance describes a human /organizational risk -trusting the model too much-rather than the model's intrinsic behavior of generating incorrect facts. Overreliance can be a consequence of this behavior, but it is not what the behavior itself is called.

In practice, you mitigate this kind of inaccuracy by grounding responses in trusted sources (for example, RAG), constraining prompts with explicit requirements, using verification steps (citations, cross-checking, tool-based validation), and adding human review for high-impact use cases.

NEW QUESTION: 19

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.

Answer Area		
Statements	Yes	No
Larger datasets can increase the cost of a generative AI solution that uses an Azure Machine Learning workspace.	☐	☐
The cost of consuming Azure OpenAI models is primarily identified by the number of input and output tokens processed.	☐	☐
The cost of custom generative AI solutions always remains the same regardless of the model version or capability used.	☐	☐

Answer:

Answer Area		
Statements	Yes	No
Larger datasets can increase the cost of a generative AI solution that uses an Azure Machine Learning workspace.	☒	☐
The cost of consuming Azure OpenAI models is primarily identified by the number of input and output tokens processed.	☒	☐
The cost of custom generative AI solutions always remains the same regardless of the model version or capability used.	☐	☒

Explanation:

Answer Area

- * Larger datasets can increase the cost of a generative AI solution that uses an Azure Machine Learning workspace. Answer: Yes
- * The cost of consuming Azure OpenAI models is primarily identified by the number of input and output tokens processed. Answer: Yes
- * The cost of custom generative AI solutions always remains the same regardless of the model version or capability used. Answer: No
- * Yes - In Azure Machine Learning, cost is driven by the resources you consume to store, process, and train. Larger datasets typically require more storage, more data transfer, and more compute time for preprocessing, training, evaluation, and experimentation. Even if you are not training foundation models, handling larger corpora can increase pipeline duration and the number/size of compute instances used, which increases overall cost.
- * Yes - For Azure OpenAI usage under Standard (on-demand), pricing is primarily tied to token-based consumption (input tokens + output tokens). The more context you send and the longer the generated responses, the more tokens you consume, and the higher the cost. This is why prompt optimization, response length controls, and grounding strategies matter for cost management.
- * No - Costs vary with model choice and capability. Different model families and versions have different price points, and larger/more capable models generally cost more per token or per unit of throughput. Additionally, architecture choices (RAG, vector search, caching), usage patterns, and throughput requirements can significantly change total cost-so it is never "always the same."

NEW QUESTION: 20

- For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Answer Area		
Statements	Yes	No
To use Microsoft 365 Copilot Chat, you must have a Microsoft 365 Copilot license.	☐	☐
Microsoft 365 Copilot Chat provides context-aware assistance in Microsoft 365 apps.	☐	☐
Microsoft 365 Copilot Chat can only access information in open files and read emails.	☐	☐

Answer:

Statements	Yes	No
To use Microsoft 365 Copilot Chat, you must have a Microsoft 365 Copilot license.	☐	☒
Microsoft 365 Copilot Chat provides context-aware assistance in Microsoft 365 apps.	☒	☐
Microsoft 365 Copilot Chat can only access information in open files and read emails.	☐	☒

Explanation:

Answer Area

- * To use Microsoft 365 Copilot Chat, you must have a Microsoft 365 Copilot license. Answer: No
- * Microsoft 365 Copilot Chat provides context-aware assistance in Microsoft 365 apps. Answer: Yes
- * Microsoft 365 Copilot Chat can only access information in open files and read emails. Answer: No
- * No - Microsoft documents and pricing guidance indicate that Copilot Chat is available for eligible Microsoft Entra ID users with a qualifying Microsoft 365 subscription , and that capabilities can vary depending on whether users also have the Microsoft 365 Copilot add-on . In other words, Copilot Chat isn't universally "Copilot-license required" in all cases; the add-on impacts which premium work features are enabled, but the baseline chat experience can be available without assigning the add-on to every user.
- * Yes - Microsoft describes Copilot Chat as a side-by-side experience in select Microsoft 365 apps (for example Word, Excel, PowerPoint, OneNote, Outlook) that is aware of the user's open content , enabling context-aware help such as summarizing, drafting, and Q & A based on what the user is working on.
- * No - The word only makes this statement inaccurate. Copilot Chat can be grounded in more than just open files and emails (for example, it can also be grounded in the web, and-depending on licensing and configuration-can use agents/connectors for broader scoped tasks). So limiting it strictly to "open files and read emails" is too restrictive compared to how Microsoft positions the feature set.

NEW QUESTION: 21

- 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
Microsoft 365 Copilot can amplify existing data governance challenges.	☐	☐
Implementing Microsoft 365 Copilot reduces data management costs.	☐	☐
Microsoft 365 Copilot can help IT teams manage data risks.	☐	☐

Answer:

Statements	Yes	No
Microsoft 365 Copilot can amplify existing data governance challenges.	☒	☐
Implementing Microsoft 365 Copilot reduces data management costs.	☐	☒
Microsoft 365 Copilot can help IT teams manage data risks.	☒	☐

Explanation:

Microsoft 365 Copilot can amplify existing data governance challenges.

The answer: Yes

Implementing Microsoft 365 Copilot reduces data management costs.

The answer: No
Microsoft 365 Copilot can help IT teams manage data risks.
The answer: Yes

Yes - Copilot relies on the permissions, sharing links, and content exposure that already exist in Microsoft 365. If an organization has oversharing (for example, broadly accessible SharePoint sites, poorly scoped Teams, unmanaged external sharing, or excessive access rights), Copilot can surface that content more easily through natural-language querying. In other words, Copilot doesn't create new permissions, but it can increase visibility of governance gaps and make the impact of weak information architecture more apparent.

No - It is not accurate to claim that implementing Copilot inherently reduces data management costs. Adoption often requires up-front investment in data hygiene, sensitivity labeling, retention, permission cleanup, DLP, and change management. Some organizations may realize productivity gains or reduced effort over time, but "reduces costs" is not a guaranteed outcome and depends heavily on the current state of governance, the scale of remediation needed, and how Copilot is rolled out.

Yes - Copilot can support IT risk management when deployed with the right controls: identity and access governance, sensitivity labels, DLP policies, retention, auditing, and compliance tooling. Because Copilot operates within the Microsoft 365 security/compliance boundary and honors existing access controls, IT can apply centralized policies to reduce leakage risk and improve overall control of how organizational data is accessed and used.

NEW QUESTION: 22

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

Answer Area

Statements	Yes	No
Lack of trust in the security and accuracy of AI solutions is a common barrier to AI adoption.	☐	☐
AI adoption efforts often fail when organizations overlook the importance of high-quality data.	☐	☐
AI adoption is most successful when implemented with business and compliance stakeholders.	☐	☐

Answer:

Answer Area

Statements	Yes	No
Lack of trust in the security and accuracy of AI solutions is a common barrier to AI adoption.	☒	☐
AI adoption efforts often fail when organizations overlook the importance of high-quality data.	☒	☐
AI adoption is most successful when implemented with business and compliance stakeholders.	☒	☐

Explanation:
The correct selections are Yes, Yes, Yes. Lack of trust is a common AI adoption barrier because users and leaders may worry about security, privacy, hallucinations, inaccurate outputs, and unclear accountability. If people do not trust the system, they will avoid using it or over-escalate decisions outside the intended workflow. AI adoption also fails when organizations overlook high-quality data.

Poor, outdated, incomplete, or poorly governed data weakens grounding, reduces output quality, and damages user confidence. Finally, adoption is strongest when business and compliance stakeholders are involved early. Business stakeholders ensure the solution solves real problems and delivers measurable value, while compliance, legal, risk, and privacy teams ensure the solution aligns with policy and regulatory obligations.

NEW QUESTION: 23

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point

Answer Area

Statements	Yes	No
For a user to access organizational data from a mobile device, the user needs a Microsoft 365 Copilot license.	☐	☐
To reason over your organizational data by using Microsoft Graph, you need a Microsoft 365 Copilot license.	☐	☐
To use the Analyst agent, you need a Microsoft 365 Copilot license.	☐	☐

Answer:

Answer Area

Statements	Yes	No
For a user to access organizational data from a mobile device, the user needs a Microsoft 365 Copilot license.	☐	☒
To reason over your organizational data by using Microsoft Graph, you need a Microsoft 365 Copilot license.	☒	☐
To use the Analyst agent, you need a Microsoft 365 Copilot license.	☒	☐

Explanation:

- * For a user to access organizational data from a mobile device, the user needs a Microsoft 365 Copilot license. No
- * To reason over your organizational data by using Microsoft Graph, you need a Microsoft 365 Copilot license. Yes
- * To use the Analyst agent, you need a Microsoft 365 Copilot license. Yes Top of Form Bottom of Form The key distinction here is between Copilot Chat capabilities available with a standard Microsoft 365 subscription and the full Microsoft 365 Copilot add-on that enables richer, in-context experiences grounded in organizational data. Mobile access to organizational data does not inherently require a Microsoft 365 Copilot license. Microsoft's Microsoft 365 Copilot app (and related mobile experiences) can provide Copilot Chat for work/school accounts with a Microsoft 365 license, so simply accessing organizational content on a mobile device is not the same as having the paid Copilot add-on. The statement claims a Copilot license is required just to access org data from mobile, which is too broad-there are mobile Microsoft 365 apps that access org data without the Copilot add-on. Reasoning over organizational data via Microsoft Graph is a core value proposition of Microsoft 365 Copilot. Microsoft documents explain that Microsoft 365 Copilot connects LLMs to your organization's content and context through Microsoft Graph and generates responses "anchored" in organizational data. That deeper integration is tied to the Microsoft 365 Copilot experience (an add-on license). Analyst is a "reasoning agent" within Microsoft 365 Copilot and Microsoft states that users with a Microsoft 365 Copilot license can use Analyst (with defined usage limits). Therefore, the Analyst agent requires the Microsoft 365 Copilot license.

NEW QUESTION: 24

For each of the following statements, select Yes if the statement is true. Otherwise, select No . NOTE: Each correct selection is worth one point.

Answer Area

Statements

Yes

No

A text-to-image generator can be used to translate content into other languages.

☐☐

A predictive analytics model can generate new marketing content for a company's online ads.

☐☐

A generative AI chatbot can engage customers in personalized conversations and recommend products.

☐☐

Answer:

Answer Area

Statements

Yes

No

A text-to-image generator can be used to translate content into other languages.

☐☒

A predictive analytics model can generate new marketing content for a company's online ads.

☐☒

A generative AI chatbot can engage customers in personalized conversations and recommend products.

☒☐

Explanation:

* A text-to-image generator can be used to translate content into other languages. Answer: No

* A predictive analytics model can generate new marketing content for a company's online ads. Answer: No

* A generative AI chatbot can engage customers in personalized conversations and recommend products.

The answer: Yes

* No - A text-to-image generator's primary function is to create images from text prompts , not translate text between languages. Translation is a natural language processing task typically handled by language models or dedicated translation services. A text-to-image model could illustrate translated content (e.g., generate an image based on a translated prompt), but it is not the tool used to perform the translation itself.

* No - Predictive analytics models are designed to predict outcomes (forecasts, probabilities, classifications) from historical patterns, such as predicting click-through rate, churn, or next-quarter demand. They are not designed to create new ad copy or marketing creatives. Generating new marketing content is a generative AI capability (text generation), not predictive analytics.

* Yes - A generative AI chatbot is well-suited to interactive, natural-language conversations . With access to product catalogs and business rules, it can ask clarifying questions, tailor responses to customer needs, and recommend products (for example, suggesting tents based on group size, season, and budget). This combines conversational generation with retrieval/recommendation logic behind the scenes, enabling personalized customer engagement at scale.

Your company manages an online catalog of office supplies. You plan to use a generative AI solution to create product descriptions for your company's website. The solution must ensure descriptions can be posted immediately after creation, enable selection/inclusion of product details, and be fast and simple for non- technical staff. What is the best type of solution to use? Select the BEST answer.

- A. custom Azure Machine Learning model
- B. the Researcher agent in Microsoft 365 Copilot
- C. an interactive AI agent
- D. a fine-tuned large language model LLM

Answer: D (LEAVE A REPLY)

The task is high-volume content generation with consistent structure and immediate publishing: product descriptions that reliably include chosen product attributes (brand, specs, materials, dimensions, use cases) and can be produced quickly by non-technical staff. The best fit is a fine-tuned LLM (D) because fine-tuning can standardize tone, format, and completeness against your catalog schema, reducing variability and minimizing manual editing before posting. With a fine-tuned model, you can strongly enforce style guidelines (length, voice, prohibited claims), and you can template prompts so staff only supply product fields and get publish-ready copy.

Option A is not best: Azure Machine Learning is excellent for predictive models but is unnecessary for straightforward text generation. B (Researcher) is optimized for multistep research across work data + web, not deterministic product copy generation. C (interactive agent) can help collect requirements, but it's more complexity than needed; the core need is consistent text generation from structured product data, which fine-tuning addresses directly while keeping user interaction simple (fill fields # generate description).

NEW QUESTION: 26

- Select the answer that correctly completes the sentence.

A _____ AI solution recognizes patterns in large and complex datasets to create new and original content.

Answer Area

A

prescriptive
predictive
generative

 AI solution recognizes patterns in large and complex datasets to create new and original content.



Answer:

Answer Area

A

prescriptive
predictive
generative

 AI solution recognizes patterns in large and complex datasets to create new and original content.



Explanation:

generative

The sentence describes an AI system that creates new, original content (for example, text, images, audio, code) by learning patterns from large datasets. That is the defining characteristic of generative AI , so "generative" is the correct completion.

To clarify the contrasts: predictive AI primarily uses historical data to forecast outcomes or estimate probabilities (for example, predicting next-quarter sales, predicting equipment failure, or classifying whether a transaction is fraudulent). It focuses on what is likely to happen or which category something belongs to , rather than producing novel content. Prescriptive AI goes a step further by recommending actions or decisions (for example, "increase inventory by X," "route tickets to team Y," "schedule maintenance next week") often using optimization and business constraints.

Generative AI, by comparison, is focused on content synthesis . It identifies statistical patterns and relationships in training data and then generates new sequences that resemble the learned distribution while being responsive to prompts. In business use cases, that translates to drafting emails and documents, summarizing reports, generating meeting notes, creating product descriptions, producing marketing copy, and building chat-based assistants. While all three categories can "recognize patterns," only generative AI is explicitly characterized by the ability to produce new and original content , which is what the question is testing.

NEW QUESTION: 27

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.)

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

Answer: A,E ([LEAVE A REPLY](#))

NEW QUESTION: 28

Your company is evaluating the use of Microsoft Copilot Studio to support business process automation and employee self-service. Which two capabilities are directly supported in Copilot Studio? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. Using agents to identify and respond to security incidents
- B. configuring document security
- C. drafting and summarizing files in Microsoft Word and PowerPoint
- D. building agents that connect to business data and automate user interactions
- E. customizing agent behavior and responses

Answer: ([SHOW ANSWER](#)**)**

Microsoft Copilot Studio is built for creating and managing custom agents that handle employee self-service and business process automation. The two capabilities that align directly to this purpose are D and E .

D is correct because Copilot Studio lets you build agents that connect to enterprise data and systems and then perform actions on behalf of users. This is the foundation for automation and self-service: the agent can answer questions using connected knowledge sources and can also trigger workflows (for example, submitting a request, creating a ticket, checking status, or updating records) through connectors and actions.

These integrations allow the agent to move beyond "chat" into real operational outcomes, which is exactly what business process automation requires.

E is correct because Copilot Studio provides the controls needed to customize how an agent behaves and responds. This includes defining conversational topics/flows, setting instructions and guardrails, shaping tone and response style, configuring fallback behavior, and controlling how generative answers are produced (for example, using approved knowledge sources). Customization ensures the agent behaves consistently with company policies and provides reliable employee experiences.

The other choices represent different Microsoft products: A is aligned to security operations tooling (Security Copilot), B relates to information protection/governance (Purview), and C is a Microsoft 365 Copilot capability inside Word/PowerPoint rather than a Copilot Studio feature.

NEW QUESTION: 29

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?

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

Answer: A ([LEAVE A REPLY](#))

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.

NEW QUESTION: 30

Your company plans to use generative AI to help build a website that will showcase various existing products. Which capability best describes a benefit of using generative AI for this project? Select the BEST answer.

- A. translating product descriptions into a different language
- B. analyzing product sales to identify patterns and trends
- C. designing a new product based on customer requirements
- D. creating product descriptions based on product specifications


Answer: D (LEAVE A REPLY)

For a product showcase website, the highest-impact, most directly relevant generative AI benefit is content creation at scale -producing consistent, high-quality product copy quickly. Option D matches a core generative AI capability: turning structured inputs (specifications such as dimensions, materials, features, compatibility, and use cases) into natural-language descriptions that are readable, persuasive, and formatted for web publishing. This accelerates catalog onboarding, reduces manual writing effort, and helps maintain a consistent tone and structure across thousands of SKUs.

Option A (translation) is also something generative AI can do, but it is a narrower, secondary capability compared to the primary website need: generating product copy from specs. Option B is predictive/analytical AI rather than generative content creation. Option C is plausible in ideation, but the scenario is about showcasing existing products, not inventing new ones. Therefore, D is the best answer because it aligns directly with the project's core workflow (product specs # publishable descriptions) and delivers measurable business value through speed, consistency, and reduced content production costs.

NEW QUESTION: 31

- Select the answer that correctly completes the sentence.
When you use Microsoft 365 Copilot connectors to connect external content to _____, your users can find, summarize, and learn from line-of-business (LOB) data by using natural language prompts.



Answer:



Explanation:
Microsoft Graph

Microsoft 365 Copilot connectors (built on Microsoft Graph connectors) are used to bring external, line-of- business content into the Microsoft 365 ecosystem by ingesting it into Microsoft Graph . Once connected, the content can be indexed and made discoverable through Microsoft Search and available for Copilot experiences, enabling users to use natural language prompts to find and summarize relevant LOB information- subject to permissions and governance controls.

The other choices don't match how Copilot connectors are positioned. Azure AI Search is an Azure indexing

/retrieval service used in custom RAG solutions, but Microsoft 365 Copilot connectors are specifically designed to surface external content through Microsoft 365 experiences via Graph. Microsoft Purview focuses on data governance, compliance, and risk management rather than being the primary ingestion target for Copilot connector content. SharePoint can store content, but the connector model is about indexing external systems into Microsoft Graph so the content becomes searchable and usable across Microsoft 365, not merely placing it into SharePoint as the destination.

So the correct completion is Microsoft Graph because that is the foundational data and indexing fabric Copilot uses to reason over organizational content with appropriate permission trimming.

Valid AB-731 Dumps shared by Actual4test.com for Helping Passing AB-731 Exam! Actual4test.com now offer the **newest AB-731 exam dumps**, the Actual4test.com AB-731 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AB-731 dumps with Test Engine here: https://www.actual4test.com/AB-731_examcollection.html (117 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 32

Your company deploys an AI-powered loan approval solution that enables applicants to request an explanation as to why their loan application was denied.

Which Microsoft responsible AI principle is this an example of?

- A. transparency
- B. fairness
- C. privacy and security
- D. inclusiveness

Answer: A (LEAVE A REPLY)

The correct answer is A. transparency. Transparency means users and stakeholders should understand when AI is being used, what the AI system can do, what its limitations are, and how important outputs or decisions were produced. In this scenario, applicants can request an explanation for why a loan application was denied. That explanation capability makes the AI-assisted decision more understandable and reviewable, which directly supports transparency. Fairness is related to avoiding unjust bias or discriminatory outcomes. Privacy and security focus on protecting personal and organizational data. Inclusiveness focuses on ensuring AI systems work for people with different abilities, backgrounds, and experiences. Because the scenario is specifically about explaining a decision, the best answer is transparency.

NEW QUESTION: 33

Your company is deploying Microsoft 365 Copilot. The deployment must provide users with access to the Researcher agent to search across data in Microsoft SharePoint. You need to recommend a licensing plan for the solution. What should you recommend?

- A. pay-as-you-go
- B. a Microsoft 365 subscription entitlement
- C. a Microsoft 365 Copilot per-user add-on license
- D. a usage-based consumption license in Azure

Answer: C (LEAVE A REPLY)

The requirement is explicit: users must have access to the Researcher agent in Microsoft 365 Copilot and use it to search across organizational content stored in SharePoint . Microsoft's licensing guidance for Researcher indicates that Researcher is available to Microsoft 365 business and enterprise users who have a Microsoft 365 Copilot add-on license (and also to certain consumer "Microsoft 365 Premium" subscriptions).

That maps directly to option C . The Researcher agent is part of the Microsoft 365 Copilot experience for work tenants; it is not enabled merely by having a baseline Microsoft 365 subscription entitlement (B).

Baseline subscriptions can provide access to Microsoft 365 apps and content repositories (like SharePoint

/OneDrive), but the Researcher agent itself is a Copilot capability that requires the Copilot add-on to unlock.

Option A (pay-as-you-go) and D (usage-based consumption license in Azure) describe consumption models that apply to Azure services or agent metering in some scenarios, but they are not the standard licensing requirement to enable the built-in Researcher agent for Microsoft 365 Copilot users. When the goal is to enable Researcher inside Microsoft 365 Copilot for staff, the practical and correct recommendation is to assign the Microsoft 365 Copilot per-user add-on license to the users who need it, ensuring their SharePoint access is already properly permissioned and governed.

NEW QUESTION: 34

Select the answer that correctly completes the sentence.

Prompt engineering is the process of _____.

Answer Area

Prompt engineering is the process of

- integrating AI-powered tools into business workflows.
- identifying and fixing errors in AI-generated content.
- designing, developing, and training generative AI models to produce contextually relevant content.
- crafting clear instructions to guide generative AI solutions in generating context-appropriate content.

Answer:

Answer Area

Prompt engineering is the process of

- integrating AI-powered tools into business workflows.
- identifying and fixing errors in AI-generated content.
- designing, developing, and training generative AI models to produce contextually relevant content.
- crafting clear instructions to guide generative AI solutions in generating context-appropriate content.

Explanation:

crafting clear instructions to guide generative AI solutions in generating context-appropriate content.

Prompt engineering is fundamentally about how you communicate intent to a generative AI model so it produces outputs that meet business expectations. The best completion is "crafting clear instructions to guide generative AI solutions in generating context-appropriate content" because it captures the practical, day-to-day discipline: shaping the input (prompt) with the right task framing, constraints, context, and output format. In real deployments, prompt engineering includes specifying the role and objective (for example, "act as a customer support agent"), providing the necessary context (product details, policy excerpts, audience), adding explicit requirements (tone, length, must/must-not statements), and defining structured output (JSON fields, bullet sections, headings). It can also include adding examples (few-shot prompting), clarifying what to do when information is missing, and instructing the model to cite only provided sources or to ask follow-up questions. These techniques reduce ambiguity, improve consistency, and lower the risk of hallucinations or off-brand responses.

The other options are not accurate definitions. "Integrating AI-powered tools into business workflows" describes solution adoption/integration, not prompt engineering. "Identifying and fixing errors in AI-generated content" is review/editing or quality assurance. "Designing, developing, and training generative AI models" is model development/ML engineering. Prompt engineering operates without changing model weights ; it's about steering model behavior through well-constructed instructions and context.

NEW QUESTION: 35

Which business requirement most closely relates to grounding a generative AI model?

- A. supporting multiple languages
- B. measuring the number of user interactions per day
- C. enabling users to interact by using natural language queries
- D. ensuring that verified company data sources are used for response generation

Answer: D (LEAVE A REPLY)

Grounding in generative AI means ensuring model outputs are based on trusted, relevant information sources rather than only on the model's general training data. In a business context, grounding is about aligning responses with verified enterprise knowledge (policies, product documentation, internal procedures, approved FAQs, etc.) so the system is more accurate, consistent, and defensible. That is exactly what option D describes: "ensuring that verified company data sources are used for response generation." In Microsoft AI solution patterns, grounding is commonly achieved using retrieval-augmented generation (RAG). With RAG, the system retrieves relevant passages from approved company repositories (for example, indexed documents or knowledge bases) and supplies them as context to the model during response generation. This reduces hallucinations, improves factual correctness, and makes answers more relevant to the organization's reality-critical when AI is used for customer support, employee helpdesks, compliance guidance, or executive reporting.

The other options do not directly address grounding. A relates to localization/multilingual capability, B is a usage/telemetry metric, and C is an interaction method (natural language interface). They can all be important requirements, but none of them ensure outputs are anchored to verified company data-the core purpose of grounding.

NEW QUESTION: 36

Your company is building a portfolio of AI-powered business solutions. Company executives want to understand how Microsoft responsible AI principles can support the company ' s long-term goals. Which benefit best demonstrates the importance of responsible AI? Select the BEST answer.

- A. guarantees that AI models provide accurate and relevant responses
- B. reduces the need for data protection policies and governance
- C. enhances stakeholder trust and fosters sustainable AI adoption throughout the organization
- D. reduces the need for executive oversight in AI decision-making

Answer: C (LEAVE A REPLY)

Responsible AI is fundamentally about earning and maintaining trust while scaling AI across the enterprise. Option C is the best answer because responsible AI practices (fairness, reliability and safety, privacy and security, transparency, accountability, and inclusiveness) reduce reputational, legal, and operational risk and make adoption sustainable over time. When stakeholders trust that AI is governed, tested, and monitored, the organization can expand AI usage confidently across business units.


The other options are incorrect because they make absolute or counterproductive claims. A is false:

responsible AI does not "guarantee" accuracy; it reduces risk and improves assurance, but no model can be guaranteed correct in all contexts. B is the opposite of reality: responsible AI increases the importance of data protection and governance; it does not reduce the need for them. D is also incorrect: responsible AI requires clear ownership and oversight, especially from leadership, because accountability is a core principle. In short, responsible AI matters because it builds stakeholder confidence and provides guardrails that support long- term, scalable, and compliant AI adoption-exactly what executives care about when investing in an AI portfolio.


NEW QUESTION: 37

Select the answer that correctly completes the sentence.

The Researcher agent in Microsoft 365 Copilot _____.



Answer:



Explanation:

uses reasoning capabilities to generate deep insights based on organizational data and the web.

The sentence is best completed by the option describing Researcher as a research-oriented reasoning agent that combines information from the web and your work data to produce deeper insights. Microsoft describes Researcher as an agent built into Microsoft 365 Copilot to tackle complex, multistep research and to help users gather, analyze, and summarize information from "the web, your work documents, or both," producing a structured output that supports decision-making. That is exactly what the completion "uses reasoning capabilities to generate deep insights based on organizational data and the web" captures.

The other dropdown options are better matches for different tools/agents: "creates visual dashboards from structured data in Excel and Power BI" is more aligned to BI/reporting workflows; "generates a pivot table and performs time series forecasting" is spreadsheet/analytics functionality; and "performs complex, multi- step, data analysis and code execution tasks over arbitrary datasets" is the hallmark positioning of the Analyst agent ("virtual data scientist") rather than Researcher. Researcher's differentiator is deep research across both organizational context and the open web, while Analyst's differentiator is data analysis and computation .

Valid AB-731 Dumps shared by Actual4test.com for Helping Passing AB-731 Exam! Actual4test.com now offer the **newest AB-731 exam dumps**, the Actual4test.com AB-731 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com AB-731 dumps with Test Engine here: https://www.actual4test.com/AB-731_examcollection.html (117 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)