

AI Transformation Leader pass4sure practice & AB-731 pdf training material



What's more, part of that itPass4sure AB-731 dumps now are free: <https://drive.google.com/open?id=1PBwtOugFUnKSAK3xq4A4hHrbeJE3rIVK>

If you free download the demos of our AB-731 study guide to have a try, then you will find that rather than solely theory-oriented, our AB-731 actual exam provides practice atmosphere when you download them, you can practice every day just like answering on the real AB-731 Practice Exam. We can help you demonstrate your personal ability and our AB-731 exam materials are the product you cannot miss.

Microsoft AB-731 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none">Identify Benefits, Capabilities, and Opportunities for Microsoft's AI Apps and Services: Focuses on mapping Microsoft's AI ecosystem including Microsoft 365 Copilot, Copilot Studio, and Azure AI Foundry Tools to real business use cases, while leveraging built-in scalability, security, and safety benefits.
Topic 2	<ul style="list-style-type: none">Identify an Implementation and Adoption Strategy for Microsoft's AI Apps and Services: Covers responsible AI principles, governance, and organizational adoption planning, including AI councils, champion programs, and an understanding of Copilot and Azure AI licensing models.
Topic 3	<ul style="list-style-type: none">Identify the Business Value of Generative AI Solutions: Covers core generative AI concepts, cost drivers, and business challenges, along with techniques like prompt engineering and RAG that enhance AI value through better data quality, security, and machine learning practices.

>> AB-731 Dumps Vce <<

Microsoft AB-731 PDF Questions - An Easy Way To Prepare For Exam

We guarantee most AB-731 exam bootcamp materials are the latest version which is edited based on first-hand information. Our educational experts will handle this information skillfully and publish high passing-rate AB-731 test preparation materials professionally. Our high quality can make you rest assured. Besides, we provide one year free updates and one year service warranty, you don't need to worry too much if how long our AB-731 Exam Guide will be valid. Once we release new version you can always download free within one year.

Microsoft AI Transformation Leader Sample Questions (Q58-Q63):

NEW QUESTION # 58

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

model?

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

Answer: D

Explanation:

Pretrained generative AI models are trained on massive, diverse datasets to gain foundational knowledge, while fine-tuned models take these pretrained weights and further train them on smaller, specific datasets to improve accuracy for narrow tasks or industries. This process aligns the model's output to specialized styles, domains, or tasks.

Key Differences and Details:

Pretrained Models (Foundational): These models (e.g., GPT-4) learn general language, concepts, and patterns from massive, broad datasets like Common Crawl. They are versatile but may lack expertise in specialized fields.

Fine-tuned Models: By adjusting the weights of a pretrained model on a smaller, labeled dataset, the model is tailored to specific applications, such as medical analysis, legal document review, or a particular brand voice.

Performance Benefits: Fine-tuning improves precision and reduces irrelevant outputs compared to a generic model.

Methodology: While pretraining is unsupervised or self-supervised, fine-tuning often uses supervised learning

Reference: <https://www.ibm.com/think/topics/fine-tuning>

NEW QUESTION # 59

Hotspot Question

Select the answer that correctly completes the sentence.

Answer Area

Microsoft 365 Copilot can be used to _____.

- monitor network traffic and alerts in real time.
- create a Microsoft Word document.
- create a list in Microsoft SharePoint.
- modify administrative permissions for Microsoft SharePoint files.

Answer:

Explanation:

Answer Area

Microsoft 365 Copilot can be used to _____.

- monitor network traffic and alerts in real time.
- create a Microsoft Word document.
- create a list in Microsoft SharePoint.
- modify administrative permissions for Microsoft SharePoint files.

Explanation:

Box: create a Microsoft Word document

Microsoft 365 Copilot can be used to _____.

Microsoft 365 Copilot can be used to create, draft, and refine Microsoft Word documents through several methods:

Draft from Scratch: You can start a new blank document and use the Draft with Copilot box (accessible via the Copilot icon or Alt + I) to enter a natural language prompt, such as "Write a sales proposal for a new product".

Reference Existing Files: You can ask Copilot to draft a new document based on up to three existing files (like a PowerPoint or another Word doc) by using the Reference a file button or typing / followed by the filename in the prompt box.

Chat-to-Document: Using the Copilot Agent in Word, you can start a project in a chat interface to ideate and then seamlessly transition that content into a structured Word document.

Template Creation: Within the Microsoft 365 Copilot app, you can select "Create" to start a document from a pre-defined template.

Reference:

<https://support.microsoft.com/en-us/office/welcome-to-copilot-in-word-2135e85f-a467-463b-b2f0-c51a46d625d1>

NEW QUESTION # 60

Hotspot Question

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

Answer:

Explanation:

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

Explanation:

Box 1: Yes

Yes - Larger datasets can increase the cost of generative AI solution that uses an Azure Machine Learning workspace. larger datasets can increase the cost of a generative AI solution using an Azure Machine Learning (AML) workspace. While the AML workspace itself has no additional charge, larger datasets increase expenses across several underlying Azure services, particularly through higher storage, compute, and data transfer requirements.

Box 2: Yes

Yes - The cost of consuming Azure OpenAI models is primarily identified by the number of input and output tokens processed. The cost of consuming Azure OpenAI models is primarily determined by the number of input and output tokens processed. This consumption-based, pay-as-you-go model calculates costs based on the total volume of text (or image/audio data) sent to the model (input) and the text generated by the model (output), usually billed per 1,000 or 1 million tokens, depending on the model and pricing page updates.

Box 3: No

No - The cost of custom generative AI solutions always remains the same regardless of the model version or capability used he cost of custom generative AI solutions varies significantly based on the model's version, complexity, and capability. While a basic chatbot might range from \$40,000 to \$150,000, advanced enterprise-grade platforms with high-risk reasoning can exceed \$400,000 to \$500,000+.

Reference:

<https://www.doit.com/blog/optimizing-ml-costs-with-azure-machine-learning>

<https://www.finout.io/blog/azure-openai-pricing>

https://medium.com/@dejanmarkovic_53716/custom-ai-solutions-cost-guide-2025-pricing-insights-revealed-cf19442261ec

NEW QUESTION # 61

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
Allowing AI models to make autonomous decisions supports the Microsoft responsible AI principle of accountability.	<input type="radio"/>	<input type="radio"/>
Regularly testing AI models for fairness and inclusiveness helps ensure they align with Microsoft's Responsible AI principles.	<input type="radio"/>	<input type="radio"/>
Protecting user data and limiting access to personal information supports the Microsoft responsible AI principles of privacy and security.	<input type="radio"/>	<input type="radio"/>

Answer:

Explanation:

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

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 # 62

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

Answer: A

Explanation:

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

