

HP

HPE0-V30


HPE AI Fundamentals

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

1. Up to Date products, reliable and verified.
2. Questions and Answers in PDF Format.



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/hpe0-v30>

Latest Version: 6.0

Question: 1

What does retrieval augmented generation do?

Response:

- A. Generates synthetic documents and labels to enlarge a small knowledge base
- B. Retrains the model on new documents whenever they are added
- C. Supplies relevant content to the model as context at query time
- D. Compresses the knowledge base so it fits inside the model's weights

Answer: C

Question: 2

A retailer indexes product photographs with a multimodal model. Broad queries such as "blue running shoe" work well, but queries naming a technical detail printed in small type on the packaging return poor results.

What best explains the difference?

Response:

- A. The index is stale and has not been rebuilt since the photographs were added
- B. Multimodal models cannot process photographs that contain any text
- C. An embedding captures overall visual meaning, not small text within the image
- D. The queries are too long for the model's context window

Answer: C

Explanation:

One vector summarizes what a picture is broadly about, so a few characters of small print contribute almost nothing to it. Recovering that detail needs text extraction indexed alongside the visual embedding.

Question: 3

What makes transfer learning from a pretrained transformer attractive for a specialized NLP task?

Response:

- A. The pretrained model already contains the specialized vocabulary of every industry
- B. General language competence transfers, reducing the labeled data required
- C. Pretraining already covers the task, so no further evaluation is required

D. A pretrained model can be fine-tuned without any labeled examples at all

Answer: B

Explanation:

The expensive part — learning how language works — has already been paid for on unlabeled text. Adapting that to a narrow task takes a fraction of the labeled examples training from scratch would demand.

Question: 4

What does an inference server contribute beyond simply loading a model?

Response:

- A. It schedules concurrent requests across the available accelerators
- B. It converts and compresses the model into a smaller format automatically
- C. It stores the conversation history for each client
- D. It retrains the model as new requests arrive

Answer: A

Explanation:

Turning a loaded model into a shared service means grouping requests to use the device efficiently while still meeting response times. Without that layer each caller would contend for the model directly.

Question: 5

What does the self-attention mechanism allow a transformer to do?

Response:

- A. Sort the input tokens by importance so that low-value tokens can be discarded early
- B. Compress the input into a single fixed-length vector before any processing begins
- C. Relate every position in a sequence to every other position when building its representation
- D. Process the sequence one element at a time so that earlier context is carried forward

Answer: C

Explanation:

Each position is represented using whatever else in the sequence is relevant to it, wherever that sits. This is the property that lets a transformer connect words far apart in a passage, which earlier sequential architectures struggled to do.

Question: 6

A RAG assistant returns a confidently wrong answer. Inspection shows the passage containing the correct answer was never among the retrieved chunks.

Where should the team focus first?

Response:

- A. The model choice, since a stronger model would infer the missing content
- B. The temperature setting, which should be reduced to zero
- C. The generation prompt, which should instruct the model more firmly
- D. The retrieval step, since answer quality is bounded by what is retrieved

Answer: D

Explanation:

If the supporting passage never reaches the model, the answer cannot be grounded no matter how the prompt is written. Chunking, the embedding model, filters and the number of results retrieved are where the fault lies.

Question: 7

A RAG system retrieves passages that are on-topic but repeatedly miss the specific detail asked about.

Which adjustment is most likely to help?

Response:

- A. Lower the temperature and shorten the prompt so the model stays closer to the context
- B. Switch the generation model for a larger one
- C. Revisit the chunking strategy
- D. Reduce the number of retrieved passages so the prompt is shorter

Answer: C

Explanation:

On-topic but imprecise results usually mean the unit of retrieval is too coarse and the ordering is too crude. Smaller units make matches specific, and re-ranking promotes the ones that actually answer the question.

Question: 8

Why is evaluating an agent harder than evaluating a single-call application?

Response:

- A. Agents cannot be run more than once on the same input
- B. Agents produce longer text, which takes more time to read
- C. Agents do not produce output that can be compared to a reference
- D. The path varies between runs, so the same input can succeed and fail for different reasons

Answer: D

Explanation:

A single call has one output to judge, whereas an agent chooses its own route and can reach a right answer badly or a wrong answer sensibly. Evaluation therefore has to look at the trajectory as well as the result.

Question: 9

A retailer wants shoppers to type a description in plain language and receive matching product photographs, with no reliance on the captions staff have written.

Which capability does this require?

Response:

- A. A text summarization model applied to the existing product captions
- B. An optical character recognition pass over every product photograph
- C. An image classifier trained on a fixed list of product categories
- D. A model that places text and images in a shared embedding space

Answer: D

Explanation:

Cross-modal retrieval works by putting a phrase and a picture into the same space so their distance is meaningful. The requirement to ignore captions is what rules out any text-only approach.

Question: 10

How does a vision transformer apply transformer machinery to an image?

Response:

- A. By processing one pixel at a time in raster order through a recurrent loop
- B. By reducing the image to a color histogram that the encoder consumes
- C. By converting the image to text with optical character recognition first
- D. By treating the image as a sequence of patches

Answer: D

Explanation:

Dividing the picture into fixed regions and ordering them produces the sequence attention needs. From there the same mechanism that relates words can relate parts of an image.

Thank You for Trying Our Product

Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>