

Grand quiz (40 marks)

Question 1:

Imagine a global healthcare company is developing two AI systems.

- The first, MedAssist, is designed to analyze medical images and detect early signs of diseases.
- The second, MedThink, is an ambitious research prototype aimed at understanding new diseases, generating hypotheses, and adapting its reasoning beyond what it has been explicitly trained for.

Questions:

- a) Based on the scenario, identify which system represents Artificial Narrow Intelligence (ANI) (weak AI) and which aspires toward Artificial General Intelligence (AGI) (strong AI). Justify your answer. (3 marks)
- b) Compare and contrast ANI and AGI in terms of their learning scope, adaptability, and decision-making autonomy, highlighting why the transition from ANI to AGI remains conceptually and technically challenging. (5 marks)
- c) Discuss one ethical or societal implication that could arise if AGI systems like MedThink were ever realized. (2 marks)

Question 2:

Scenario:

A global tech company launches *LexiGPT*, a powerful Large Language Model (LLM) designed to assist customer service teams by answering complex user queries, generating reports, and even drafting policy documents. Over time, users start noticing that while *LexiGPT* produces fluent and confident responses, some of its answers contain subtle factual errors or reflect biased viewpoints based on the data it was trained on.

Questions:

- a) Explain the underlying working principle of LLMs, focusing on how they generate contextually relevant text. (3 marks)
- b) Discuss how LLMs differ from traditional NLP in their learning approach and generalization capabilities. (3 marks)
- c) Analyze the key challenges and risks of relying on LLMs like *LexiGPT* for decision-making in professional or academic settings. (2 marks)
- d) Propose one practical solution or safeguard that can reduce the negative impacts of misinformation or bias in LLM-generated content. (2 marks)

Question 3:

Scenario:

You are part of an AI research team designing a delivery-drone navigation system. The drone must find the most efficient path between warehouses and delivery points while avoiding obstacles such as buildings and restricted air zones.

The system currently uses a Breadth-First Search (BFS) algorithm, but the team is debating whether to replace it with A* or Greedy Best-First Search to improve performance and efficiency.

Questions:

- a) Explain how uninformed and informed search algorithms differ in terms of the information they use to explore a state space. (3 marks)
- b) Compare BFS, Greedy Best-First Search, and A* in terms of their optimality, completeness, and use of heuristics, applying your answer to the drone scenario. (4 marks)
- c) Conceptually discuss the role of a heuristic function in improving the efficiency of informed search algorithms, and explain what could go wrong if the heuristic is poorly designed. (3 marks)

Question 4

Scenario:

A startup called VisionMed is developing an AI-assisted diagnostic system that analyzes chest X-rays to detect early signs of lung diseases.

The system uses Convolutional Neural Networks (CNNs) to classify images as “Healthy” or “Abnormal.” However, the team notices inconsistent accuracy — sometimes the model performs well on training data but fails to generalize to new, unseen images.

Questions:

- a) Explain how CNNs extract visual features from medical images through convolution, pooling, and activation functions. (3 marks)
- b) Discuss how model architecture choices (e.g., depth, number of layers, kernel size) can affect both the accuracy and computational efficiency of the system. (3 marks)
- c) Identify one likely cause of the inconsistent model performance mentioned in the scenario, and suggest a regularization or optimization technique (e.g., batch normalization, dropout) to address it. Justify your choice conceptually. (2 marks)
- d) Evaluate one ethical or practical concern associated with deploying CNN-based medical vision systems in real-world clinical settings. (2 marks)