

Do not start attempting the test paper until you are asked to do so.



ENROLLMENT NUMBER:

STUDENT NAME:



Maximum Marks : 100
Duration : 40 Minutes
Total Questions : 35

Test Paper Type

CC

1st
iAIO INTERNATIONAL
ARTIFICIAL INTELLIGENCE
OLYMPIAD
2025-2026

Instructions for the Candidates

1. For filling up information about yourself on the OMR sheet, you will be given additional ten minutes before the start of the examination.
2. Write your 12-digit iAIO Enrollment Number and your name on top of the Question Paper in the given space.
3. Do not forget to sign the OMR sheet. Also, write your Roll No. on the Question Paper Booklet, and do not take this home.
4. The Question Paper Booklet consists of 35 questions, divided into two sections.
5. Section-A: Artificial Intelligence (30 questions); Section-B: Scholar's Zone (5 questions).
6. Each question of Section-A carries 2.5 marks, and that of the Scholar's Zone carries 5 marks.
7. All the questions are compulsory, and there is no negative marking.
8. Use of a calculator or any other devices in the examination is strictly prohibited.
9. Choose only ONE OPTION as an answer.
10. BLUE/BLACK ball pen is preferred to darken the circle; however, in case of non-availability, an HB pencil can be used. Mark your choice of answer in the OMR sheet by darkening a circle, as shown below.



Do not start attempting the test paper until you are asked to do so.

ENROLLMENT NUMBER:

STUDENT NAME:

Maximum Marks : 100
Duration : 40 Minutes
Total Questions : 35

Test Paper Type

CC

Class **8**

1st
iAIO INTERNATIONAL
ARTIFICIAL INTELLIGENCE
OLYMPIAD
2025-2026

Instructions for the Candidates

1. For filling up information about yourself on the OMR sheet, you will be given additional ten minutes before the start of the examination.
2. Write your 12-digit iAIO Enrollment Number and your name on top of the Question Paper in the given space.
3. Do not forget to sign the OMR sheet. Also, write your Roll No. on the Question Paper Booklet, and do not take this home.
4. The Question Paper Booklet consists of 35 questions, divided into two sections.
5. Section-A: Artificial Intelligence (30 questions); Section-B: Scholar's Zone (5 questions).
6. Each question of Section-A carries 2.5 marks, and that of the Scholar's Zone carries 5 marks.
7. All the questions are compulsory, and there is no negative marking.
8. Use of a calculator or any other devices in the examination is strictly prohibited.
9. Choose only ONE OPTION as an answer.
10. BLUE/BLACK ball pen is preferred to darken the circle; however, in case of non-availability, an HB pencil can be used. Mark your choice of answer in the OMR sheet by darkening a circle, as shown below.



Section-A (Artificial Intelligence)

- What mathematical operation is performed on the output of a neuron before it becomes the input to the next layer in a Deep Learning network?
 - Feature Scaling
 - Weight Initialisation
 - Activation Function
 - Loss Calculation
- An AI student wants to use Python to build a simple neural network. Which core Python library is primarily used to handle the large arrays and high-level mathematical operations that form the weights and data structures of the network?
 - Pandas
 - Scikit-learn
 - NumPy
 - Matplotlib
- Which of the following tasks is not a typical application of Deep Learning (DL) models?
 - Real-time translation of spoken language.
 - Autonomous navigation of self-driving vehicles.
 - Rule-based sorting of network packets by IP address.
 - Generating novel, high-resolution images from a text description.
- Why is the non-linearity introduced by Activation Functions (like ReLU or Sigmoid) essential in a Deep Neural Network?
 - To guarantee the network will converge to the correct solution quickly.
 - To ensure that the model's final prediction is always a probability between 0 and 1.
 - To enable the network to learn complex relationships in the data.
 - To prevent the data values from becoming too large during the forward pass.
- A self-driving car's Deep Learning system receives live video from the front camera. This raw pixel information is first processed by the first set of mathematical operations. Which part of the artificial neural network architecture is performing this initial data reception and transformation?
 - The Output Layer, transforming the image into a decision.
 - The Input Layer, preparing the raw data for processing.
 - The first Hidden Layer, detecting simple edges and lines.
 - The Optimiser, adjusting the layer's weights based on the pixels.
- The Keras library in Python is often described as a high-level API. What does 'high-level' primarily refer to in this context?
 - It can only be used by advanced programmers.
 - It provides simple tools that hide complex, mathematical code.
 - It requires extremely high-end computer hardware to run efficiently.
 - It automatically manages memory allocation for the user.
- Which sequence correctly outlines the necessary steps to build and train a simple Deep Learning model using a Python library like Keras?
 - Define the network structure (add layers).
 - Evaluate the trained model on test data.
 - Import the necessary library (e.g., Sequential, Dense).
 - Compile the model (specify optimiser and loss function).
 - Train the model (fit the data).
 - 3 → 1 → 4 → 5 → 2
 - 1 → 3 → 4 → 5 → 2
 - 3 → 4 → 1 → 5 → 2
 - 1 → 4 → 3 → 5 → 2
- Which Python library function would a data scientist use to split their labelled image dataset into separate sets for training, validation, and final testing?
 - pd.read_csv() from Pandas.
 - train_test_split() from Scikit-learn.
 - np.concatenate() from NumPy.
 - model.add() from Keras.

9. Consider a situation where a model's accuracy on the training data is 99%, but its accuracy on the validation data is only 55%. What should the programmer's first logical step be?
- A. Add more layers to the model to increase its complexity.
 - B. Increase the number of training epochs to 100 or more, exposing model to all data repeatedly.
 - C. Reduce model complexity or use regularisation to counter overfitting.
 - D. Switch the loss function from 'mse' to 'binary_crossentropy'.
-
10. The Optimiser (e.g., 'Adam' or 'SGD') is defined in the model.compile() step. What is its core function during the training process?
- A. It calculates the numerical value of the loss (error).
 - B. It determines the probability of the correct class at the output layer.
 - C. It adjusts the model's weights to minimise the calculated loss.
 - D. It organises the layers of the model into the correct sequence.
-
11. An Advanced Computer Vision model is known to be primarily reliant on Gradient Calculation and Non-Maxima Suppression. What is the likely function of this specific model?
- A. Identifying the speed and direction of a moving object.
 - B. Clearly finding the boundaries and outlines of objects within a scene.
 - C. Generating completely new, photo-realistic images from text prompts.
 - D. Translating written text found in an image into a digital document.
-
12. Which Computer Vision technique involves not just drawing a simple bounding box around an object, but classifying every individual pixel in an image as belonging to a specific object or to the background?
- A. Instance Segmentation
 - B. Object Tracking
 - C. Motion Estimation
 - D. Feature Matching
-
13. A Computer Vision system is designed to use a camera to overlay virtual furniture onto a user's live video feed of their living room, allowing them to see how a new sofa would look in the space.



This is a real-world application of:

- A. Image Classification
- B. Generative Adversarial Networks (GANs)
- C. Augmented Reality (AR)
- D. Virtual Reality (VR)

-
14. Which characteristic of the YOLO (You Only Look Once) family of Object Detection algorithms makes them revolutionary for real-time applications like drone navigation?
- A. Their use of Instance Segmentation, which is pixel-perfect and provides detailed object boundaries.
 - B. Their exceptionally high speed from performing detection in a single network pass.
 - C. Their exclusive focus on detecting only static, non-moving objects.
 - D. Their complete reliance on feature matching instead of neural networks.
-
15. In Computer Vision, a filter (or kernel) is a small matrix of numbers that is passed over the input image to modify it. Which type of filter is specifically used to soften image detail and reduce noise?
- A. Edge Detection Filter
 - B. Sharpening Filter
 - C. Blurring Filter
 - D. Feature Extraction Filter
-
16. An AI model is designed to detect a rare disease (Positive case) in patients. In the results, the model identifies a patient as having the disease when they are, in fact, healthy. This specific error is known as a:
- A. True Positive (TP)
 - B. True Negative (TN)
 - C. False Positive (FP)
 - D. False Negative (FN)
-
17. Why is the concept of generalisation the ultimate test for an AI model's success in the real world?
- A. Generalisation means the model can perfectly memorise the training data.
 - B. Generalisation measures the model's ability to perform well on unseen, real-world data.
 - C. Generalisation ensures the model runs extremely fast on a high-speed processor.
 - D. Generalisation is only relevant for linear regression models, due to their simple structure.
-
18. If a classification model has a high F1 Score and is deployed, this suggests that the model is successfully balancing the need to:
- A. Only train quickly and not worry about accuracy.
 - B. Minimise both False Positives (FP) and False Negatives (FN).
 - C. Maximise True Positives (TP) and minimise True Negatives (TN).
 - D. Ensure the model is resistant to overfitting on the training data.
-
19. What is the name of the modern family of Generative AI models that starts with an image of random static noise and refines it over many steps to produce a coherent image from a text prompt?
- A. Generative Adversarial Networks (GANs)
 - B. Diffusion Models
 - C. Recurrent Neural Networks (RNNs)
 - D. Rule-based Expert Systems
-
20. Which digital standard allows computers and electronic instruments to communicate musical information (such as note pitch, duration, and volume) and is often used as the structured input for AI music generation systems?
- A. MP3
 - B. WAV
 - C. MIDI
 - D. JPEG
-
21. In the process of Prompt-to-Image Generation, why is the model mapping the text prompt into a "latent space" a crucial step?
- A. The latent space is where the final image is stored and exported.
 - B. The latent space is a compressed conceptual map used for idea blending.
 - C. The latent space is where the AI checks its database for copyright infringement.
 - D. The latent space is simply the final output layer of the network.
-
22. A video game company wants its Non-Player Characters (NPCs) to engage in more fluid, context-aware conversations with the player that are not pre-scripted. Which type of AI model is essential for giving the NPCs this conversational ability?
- A. A Generative Adversarial Network (GAN) used for image creation.
 - B. A Large Language Model (LLM) using Transformer architecture.
 - C. A simple Machine Learning Classifier trained on structured data.
 - D. An Advanced Computer Vision Model focused on visual tracking.

23. A city uses an AI system with sensors in its public recycling bins. The AI analyses fill levels and optimises garbage truck routes in real-time. This application primarily contributes to Smart City objectives by improving:
- A. The city's general population forecast and demographic planning.
 - B. Operational efficiency by reducing fuel consumption and vehicle wear and tear.
 - C. The accuracy of the city's traffic congestion model and predictive road maintenance.
 - D. The public's willingness to recycle through automated feedback and reward systems.

24. Consider the application of AI in promoting Sustainability:


Statement 1: AI can analyse satellite imagery to guide precision farming, reducing water and pesticide use.

Statement 2: AI can predict where illegal logging or deforestation is most likely to occur, helping preventative efforts.

Statement 3: AI can manage waste by classifying and sorting materials, but it cannot currently optimise recycling routes.

Which of the statements are True?

- A. Statement 1 and 2 only
 - B. Statement 1 and 3 only
 - C. Statement 2 and 3 only
 - D. Statement 1, 2, and 3
25. What is the term for the AI application in medicine that uses sensors and wearable devices to constantly track a patient's vital signs and automatically alert a doctor if conditions become dangerous?



- A. Personalised Medicine
- B. Early Diagnosis
- C. Remote Monitoring
- D. Surgical Automation

-
26. The core philosophy of “AI for a Better World” focuses on applying Artificial Intelligence to address which category of global issues?
- A. Optimising global supply chains and international trade to foster economic growth.
 - B. Address critical humanitarian challenges like climate change, and disaster response.
 - C. Streamlining public services and government administration to improve national efficiency.
 - D. Accelerating fundamental scientific discovery in fields like advanced mathematics and physics.
-
27. Why is the ethical principle of Informed Consent more complex in the age of Big Data and AI?
- A. Because it’s impossible to know future data use by AIs not yet built.
 - B. Because AI models are computationally too complex to train without data.
 - C. Because all data is now automatically deleted after 24 hours.
 - D. Because the speed of AI development is too slow to ask for consent.
-
28. A social media platform uses an AI to moderate content. To ensure Fairness, the system is constantly tested on posts written in English, Spanish, and Mandarin. This specific act is a strategy to prevent:
- A. Data Encryption Failure
 - B. Representation Bias
 - C. Training Overfitting
 - D. Logical Puzzle Flaws
-
29. The term “Black Box AI” is used to describe a system whose internal decision-making process is opaque or not understandable. Why does the complexity of Deep Neural Networks often lead to this “Black Box” issue?
- A. Because the underlying Python code is fundamentally unreadable and lacks proper documentation.
 - B. Their massive number of weights and complex non-linear relationships defy human interpretation.
 - C. Because all training data is always automatically encrypted by default to protect user privacy.
 - D. Because they inherently suffer from severe overfitting, causing unstable and unpredictable decisions.
-
30. A loan approval AI system is found to disproportionately reject applicants from a specific, low-income neighbourhood, even when their credit scores are comparable to approved applicants from wealthier areas. This bias is most likely due to:
- A. Historical Bias, reflecting human discrimination in the training data.
 - B. Evaluation Bias, where the test set was too small to be meaningful.
 - C. Group Bias, where the developers only included people with high credit scores.
 - D. Data Security issues, where the data was leaked.

31. Consider the following statements about Deep Neural Networks (DNNs):

Statement 1: The depth of a network (having multiple hidden layers) allows the model to map simple features into highly complex, abstract concepts.

Statement 2: DNNs require significantly less training data than traditional machine learning algorithms to achieve comparable performance.

Statement 3: In a well-trained network, the error of the model is calculated in the output layer and then propagated backwards to update the weights of the hidden layers.

Which of the statements are True?

- A. Statement 1 and 2 only
 B. Statement 1 and 3 only
 C. Statement 2 and 3 only
 D. Statement 1, 2, and 3

32. A student writes the following Keras code snippet to define the output layer for an image classification task with 5 distinct categories:

```
# The input image is 28x28 grayscale
model.add(Dense(5, activation='relu'))
```

What is the logical error in the student's reasoning for this layer?

- A. The ReLU activation function is incorrect for a multi-class final layer.
 B. The Dense layer is completely unnecessary for an output layer.
 C. The code lacks an Optimiser call, which must be part of the final layer.
 D. The number of neurons should be 6 (5 classes + 1 for 'none').

33. If a Convolutional Neural Network (CNN) in a Computer Vision task is failing to recognise objects that are rotated or viewed at a sharp angle, what feature of the training data is most likely missing, and what step should the developers take?

- A. Missing: Labelled bounding boxes.
 Step: Use a different optimiser like 'Adam' to improve convergence speed.
 B. Missing: Data from different light conditions.
 Step: Switch the output function to Sigmoid.
 C. Missing: Geometrically diverse examples.
 Step: Use Data Augmentation to fix the rotational invariance issue.
 D. Missing: High-resolution images.
 Step: Use a shallower network.

34. An AI system for detecting forest fires yielded the following results over a season:

- True Positives (TP): 45 (Correctly identified fires)
- False Positives (FP): 5 (Falsely identified smoke as fire)
- False Negatives (FN): 5 (Missed actual fires)
- True Negatives (TN): 445 (Correctly identified non-fire days)

What is the calculated Recall (Sensitivity) of the system?

- A. 90% B. 91.8% C. 88.2% D. 95%

35. The following question presents an assertion followed by a reason. Read carefully and choose the correct answer.

Assertion (A): *In high-stakes AI applications like medical diagnosis or criminal justice risk assessment, Explainable AI (XAI) is mandatory.*

Reason (R): *Without XAI, human experts cannot audit the system, verify its logic, or be held accountable for life-altering decisions made by the machine.*

- A. Both A and R are true, and R is the correct explanation for A.
- B. Both A and R are true, but R is not the correct explanation for A.
- C. A is true, but R is false.
- D. A is false, but R is true.

