

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

iAIO ^{1st} 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)

1. A mobile application suggests the best route home by predicting traffic using real-time data and past traffic patterns. This scenario primarily demonstrates which key concept of AI?
 - A. Data Collection
 - B. Machine Learning
 - C. Computer Vision
 - D. Data Visualisation

2. A student claims: “An algorithm is only useful if it uses complicated mathematical symbols; simple step-by-step recipes cannot be considered algorithms in AI.” What is the flaw in their reasoning?
 - A. The flaw is that an algorithm must always be written in a specific coding language, like Python, to be considered a formal, abstract, universal set of instructions.
 - B. The flaw is that algorithms must be complex enough to execute on multiple cores (parallel processing) to qualify as AI instructions.
 - C. The flaw is that an algorithm is fundamentally a finite set of unambiguous, logical steps, regardless of the complexity of the steps themselves.
 - D. The flaw is that algorithms are simply the final output data produced by a trained model, not the input instructions.

3. Why is diverse and accurate data crucial for building fair and effective AI systems?
 - A. Because collecting the largest amount of data possible is the only way to guarantee the AI makes smart decisions.
 - B. Because the AI needs to learn from a wide variety of examples to ensure its conclusions are accurate and unbiased for all groups of people.
 - C. Because using very clean, organised data helps the AI train much faster and with fewer technical errors.
 - D. Because if the AI makes even one mistake on the training data, it will not be able to learn anything at all.

4. An AI model used for hiring job candidates was trained on data from a company where, historically, most managers were men. The AI now tends to recommend more male candidates for manager roles. This is an example of:
 - A. AI Efficiency
 - B. AI Accuracy
 - C. AI Bias
 - D. AI Security

5. Which of the following data types would be the most appropriate choice for storing the exact mass of an apple in grams (e.g., 155.3 g)?
 - A. Integer (int)
 - B. String (str)
 - C. Float (float)
 - D. Boolean (bool)

6. In Python, what is the significance of indentation (the white space before a line of code) inside an if statement or a for loop?
 - A. It is only for making the code look tidy and visually understandable.
 - B. It defines a block of code, telling the interpreter which lines belong to the control structure.
 - C. It allows Python to dynamically adjust the speed of the code execution based on the loop’s size, enabling efficient processing.
 - D. It must always be exactly two spaces wide for code validation.

7. The computer runs an if statement block. Arrange the steps in the order the Python interpreter executes the code.
 1. The program executes the indented block of code.
 2. The interpreter checks the truth of the condition (e.g., Is $x > 10$?).
 3. The program moves on to the next line of code outside the if block.
 4. If the condition is True, the indented block is selected for execution.
 - A. 2→4→1→3 B. 4→1→3→2
 - C. 2→1→4→3 D. 1→4→3→2

8. The following question presents an assertion followed by a reason. Read carefully and choose the correct answer.
- Assertion (A): Adding comments (lines starting with #) to a Python program is essential for the computer to understand the code logic before execution.
- Reason (R): Comments are lines of text intended only for human readers, and the Python interpreter completely ignores them during the execution process.
- A. Both A and R are true, and R is the correct explanation of A.
B. Both A and R are true, but R is NOT the correct explanation of A.
C. A is true, but R is false.
D. A is false, but R is true.


9. What is the syntax error in this Python code snippet?

```
if age > 18
    print("Allowed")
print("Done checking")
```

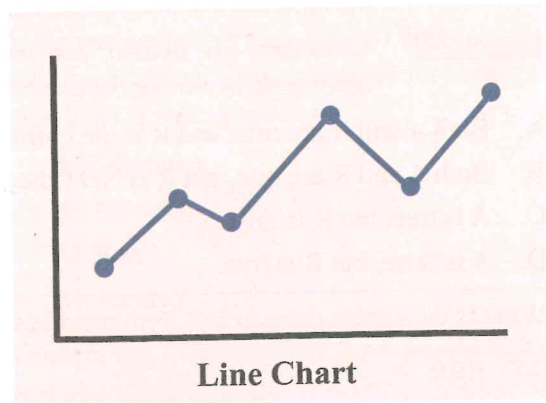
- A. The lack of parentheses (()) around the condition, making it if (age > 18).
B. The use of the comparison operator (>) instead of the assignment operator (=).
C. The missing colon (:) at the end of the if statement line.
D. The print statements are not enclosed in a function.
10. Consider the following statements about the preparation of data for Machine Learning:
- Statement 1: All data must be cleaned to remove errors or missing values.
Statement 2: Data should be split into training, validation, and testing sets.
Statement 3: All numeric data must be converted into text strings.
- Which of the statements are TRUE?
- A. 1 and 2 only B. 1 and 3 only C. 2 and 3 only D. 1, 2, and 3
11. Imagine a sorting robot on a conveyor belt that uses its camera (AI) to look at a piece of fruit and must instantly decide if the fruit is a "Banana," an "Apple," or an "Orange." What specific type of fundamental Machine Learning task is this robot performing?
- A. Regression B. Classification
C. Data Visualisation D. Predictive Maintenance
12. A marketing company wants to use ML to predict the exact amount of money a customer will spend in the next month based on their past purchase history. Which specific type of ML problem should they set up?
- A. Classification, because the output is a type of customer.
B. Regression, because the output is a precise, continuous numerical value (money).
C. Unsupervised Learning, because the customers are not labelled.
D. Clustering, because the customers are grouped by age, revealing important segmentation insights.
13. An ML model predicts students' final exam scores. For one student, the model predicts a score of 95, but the actual score is 60. This large difference is known as Error or Loss. Why is reducing this Error a critical step in training the model?
- A. To ensure the model runs faster on the computer.
B. To prove that the training data was perfectly clean and completely unbiased.
C. To make the model's predictions more accurate and reliable.
D. To prevent the model from becoming too simple to recognise complex patterns.

14. You have collected the daily rainfall for your city over the last month and want to display the fluctuation in rainfall over this time period. Which type of graph is most appropriate?

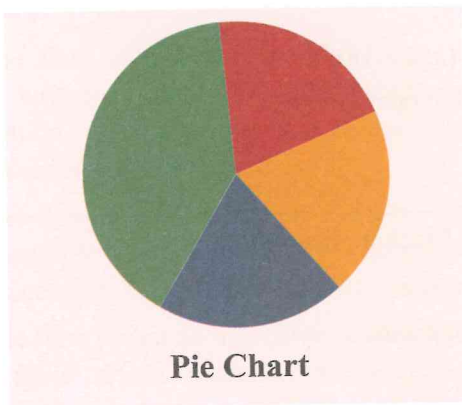
A.




B.



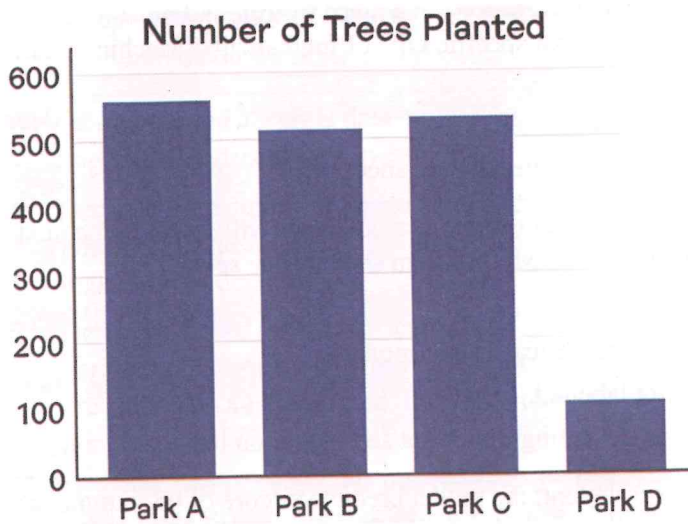
C.



D.



15. Consider the given chart, which shows the number of trees planted in four parks. Which park's data point is the clearest example of an outlier?



- A. Park A
- B. Park B
- C. Park C
- D. Park D

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

Assertion (A): A Pie Chart is the most appropriate visualisation to compare the total sales revenue generated by four different product lines (A, B, C, and D).

Reason (R): A Pie Chart's primary purpose is to show the proportion or percentage that each category contributes to the whole (the total revenue).

- A. Both A and R are true, and R is the correct explanation of A.
 B. Both A and R are true, but R is NOT the correct explanation of A.
 C. A is true, but R is false.
 D. A is false, but R is true.
-
17. Why is data visualisation considered essential when working with large datasets in AI?
- A. It reduces the total amount of data that the computer needs to store on the hard drive.
 B. It helps humans quickly identify patterns, trends, and outliers in the data.
 C. It is the first step in creating the final prediction made by the AI model.
 D. It guarantees the data is perfectly balanced and ready for the machine learning process.
-
18. Which Python library is most commonly used to create visualisations like bar charts and line graphs?
- A. NumPy B. Pandas
 C. Matplotlib D. TensorFlow
-
19. How do neural networks learn to recognise complex patterns, like faces or handwriting?
- A. By following a detailed step-by-step instruction list that the programmer wrote for every possible pattern.
 B. By adjusting the strength of the connections (weights) between artificial neurons based on training data.
 C. By simply memorising a giant library of all existing photos and looking up a match instantly.
 D. By deleting the least important layers and simplifying the network's structure after every mistake.




20. Match the components of the human brain to the elements of an artificial neural network.

Component (Human Brain)		Artificial Counterpart (Neural Network)	
1.	Brain Cell	P.	Input Layer
2.	Sense Organs	Q.	Weight/Connection
3.	Synapse	R.	Neuron

- A. 1-P, 2-Q, 3-R
 B. 1-R, 2-P, 3-Q
 C. 1-Q, 2-R, 3-P
 D. 1-R, 2-Q, 3-P
-
21. Why is a neural network often described as a black box in some contexts?
- A. Because the internal data is encrypted to protect it from being stolen by other programs.
 B. Because it is challenging to understand or trace the exact steps and millions of weight adjustments that lead to the final decision.
 C. Because its design is always inspired by human brain anatomy, which is too complex to diagram fully and understand conceptually.
 D. Because the model takes too long to train, making the full process economically opaque.
-
22. If a neural network is designed to predict the price of a house (a Regression problem), what kind of value would the Output Layer produce?
- A. A predefined category label (e.g., 'Luxury' or 'Standard'), representing a single classification.
 B. A single continuous numerical value, like 355,000.00.
 C. A list of probabilities indicating the certainty of the prediction.
 D. A list of all comparable house features (e.g., size, location) used in the prediction.
-
23. A simple neural network has 5 inputs and 3 outputs. If it has only one hidden layer with 8 neurons, how many total connections (weights) are there between the Input Layer and the Hidden Layer?
- A. 5 B. 16
 C. 40 D. 24

24. A robot arm is learning to stack a block. If the block falls (error), a signal travels to the beginning through the arm's joints and wires to immediately adjust the power and angle settings for the next try. What is this crucial step of sending the error signal through the system to adjust the internal settings called in a neural network?
- A. Forward Propagation B. Gradient Descent
C. Backpropagation D. Loss Minimisation

25. Match the IoT device with its primary function..075

IoT Device	Function
1.  Smart Thermostat	a. Monitors a home for motion and sends alerts.
2.  Smart Security Camera	b. Learns your heating habits and adjusts temperature.
3.  Smart Light	c. Turns on and off based on light levels or movement.

- A. 1-a, 2-b, 3-c B. 1-b, 2-a, 3-c
C. 1-c, 2-b, 3-a D. 1-a, 2-c, 3-b

26. A smart bin (an AI-IoT device) sends an alert to the council, stating it is full, when in reality, it is only half-full. What is the most responsible and effective first step the council can take to permanently fix this recurring error?
- A. Immediately reset the bin's sensor to its default factory settings and resume normal service immediately.
B. Dismiss the alert, as inaccurate data is expected in an initial deployment phase and wait for repeated errors.
C. Analyse sensor data to identify the cause of the prediction error, and use that correct information to re-train the AI model.
D. Deploy a secondary monitoring system using Computer Vision to verify the bin's status and confirm the reading is wrong.

27. The image shows a factory worker monitoring machines. How does Artificial Intelligence (AI), integrated with IoT sensors, provide the greatest value in preventing machine breakdowns in this industrial setting?



- A. By displaying a real-time dashboard with the current temperature and vibration levels for every machine.
B. By sending an alert to an engineer whenever a machine operates above a pre-set safety limit.
C. By analysing continuous sensor data to predict the optimal time for maintenance before a failure occurs.
D. By creating a historical report of which machines have required the most repairs in the past year.

28. Consider the following statements about Ethical AI Development:

Statement 1: Diverse teams, not just coders, should be involved in the design process to catch varied forms of bias.

Statement 2: AI systems should be regularly checked (audited) for potential bias or unfair outcomes.

Statement 3: AI should only collect the minimum amount of user data required for the task.

Which of the statements are TRUE?

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

29. An AI-generated personal news feed only shows you articles and videos that align with your existing beliefs. What is the long-term ethical risk of this behaviour concerning social discourse and critical thinking?
- It makes the model overfit to your profile, causing its predictions to fail completely when you encounter new topics.
 - It limits your worldview and reinforces existing biases by creating a 'Filter Bubble'.
 - It will cause the AI system to confuse news articles with general search queries, requiring continuous maintenance.
 - It significantly reduces the latency (time delay) for loading articles, making your overall internet experience faster.

30. An AI is trained using a biased dataset that only includes images of light-skinned hands when classifying medical scans. When used for all people, the AI incorrectly diagnoses dark-skinned people. This is an example of what kind of ethical failure?
- Transparency
 - Accountability
 - Lack of Fairness
 - Privacy

Section-B (Scholar's Zone)

31. A student writes the following Python code snippet: The student asserts that this code will print the sequence of numbers 1, 2, 3, 4, 5. What is the precise flaw in the student's reasoning regarding Python's range() function?

```
for i in range(5):
    print(i)
```

- The range(N) function generates numbers including N itself, leading to an off-by-one error.
 - The range(N) function, when given a single argument N, generates numbers in descending order from N-1 down to 0.
 - The range(N) function starts from 0 and stops before N.
 - The loop variable i must be explicitly initialised to 1 before the for loop to achieve the desired starting value.
32. Arrange the steps in a typical AI-IoT interaction sequence for a smart refrigerator that detects when milk is running low:
- The AI decides milk is running low based on weight sensor data.
 - The smart refrigerator's weight sensor (IoT component) measures the milk level.

- The refrigerator connects to the internet and sends an order to a grocery app.
 - The weight sensor sends the raw data to the refrigerator's internal processor.
- | | |
|---------------|---------------|
| A. 2, 1, 3, 4 | B. 4, 1, 3, 2 |
| C. 2, 4, 1, 3 | D. 3, 2, 4, 1 |

33. A Python program executes the following lines, and the output shown is 7.

```
a = 3
b = 4
c = a + b
print(c)
```

- To achieve the output 34 (as a string) while keeping the initial variable declarations (a = 3 and b = 4) unchanged, which line requires the essential modification?
- The line print(c) must be changed to print(a + b)
 - The line c = a + b must be changed to c = str(a) + str(b)
 - The line c = a + b must be changed to c = a * 10 + b
 - a = 3 must be changed to a = "4" and b = 4 to b = "3"


34. Which statement(s) best describe the Activation Function in a single neuron?

Statement 1: It decides if the neuron should be "fired" (activated) and what output value it should pass forward.

Statement 2: It calculates the initial input value received from the previous layer. Statement 3: It helps introduce non-linearity, allowing the network to learn complex, non-straight-line patterns.

- A. Statement 2 only
B. Statement 1 and Statement 2 only
C. Statement 1 and Statement 3 only
D. All statements are true.

35. A small robotic dog is placed in a virtual maze. The robot must learn how to reach the exit on its own through trial-and-error. It is never shown the correct path beforehand. The system gives the robot a +5 point reward when it moves toward the exit and a -2 point penalty when it bumps into a wall. The robot's goal is to find the fastest path to maximize its total score.



Which component of this Reinforcement Learning system is the +5 point reward primarily intended to influence?

- A. The design of the Environment, by constantly and randomly changing the location of the maze walls every time the robot hits one.
B. The Agent's Policy, by strengthening the likelihood that the robot will repeat the successful action in the future.
C. The initial Training Set of labelled images, by adding new pictures of the maze to the database.
D. The total Computational Power used, by forcing the computer to run faster calculations.

