

# SCO INTERNATIONAL

## ARTIFICIAL INTELLIGENCE OLYMPIAD

### CLASS 9 SAMPLE QUESTION PAPER

A comprehensive document for schools, teachers, parents, and students

**Designed from Class 9 AI pathways and aligned with SCO's guided preparation, practice, reporting, and future-ready AI literacy.**

- age-fit AI, machine learning, NLP, Python, ethics, and application-based reasoning for Class 9 learners globally
- professionally formatted question paper with compact question-number tags, answer key, and explanations
- ready for school, teacher, parent, and student use as a website/PDF learning resource

AI	Machine Learning	NLP	Python	Applications
Data	Logic	Fairness	Capstone	Practice

### Paper Overview

Exam Name	SCO International Artificial Intelligence Olympiad
Class / Grade	Class 9 / Grade 9
Question Paper Set	Set S
Total Questions	50 multiple-choice questions
Duration	60 minutes
Question Type	Objective type with one correct answer
Sections	General Questions and Achievers Section
Use	Website sample paper, classroom practice, and guided preparation

## Candidate Instructions

- Read every question carefully before marking the answer.
- There is only one correct answer for each multiple-choice question.
- Calculators, mobile phones, and digital aids are not allowed unless specifically permitted by the examination authority.
- Use the question paper for rough reading only; answers must be marked as instructed by the invigilator or online system.
- Passages, code snippets, and case details are part of the question block and should be read fully before choosing an option.
- This document includes the answer key and explanations for guided learning and website sample-paper use.

## Section A: General Questions

**Q.1** A convolutional neural network (CNN) has the following layers:

Input layer:  $128 \times 128 \times 3$

Convolution layer: 32 filters, kernel size =  $3 \times 3$ , stride = 1, padding = "same"

MaxPooling layer: Pool size =  $2 \times 2$ , stride = 2

What is the output size of the MaxPooling layer?

- A.  $64 \times 64 \times 32$
- B.  $126 \times 126 \times 32$
- C.  $128 \times 128 \times 32$
- D.  $62 \times 62 \times 32$

**Answer: A**

**Explanation:** With padding "same", the convolution output remains  $128 \times 128$  and 32 filters create 32 channels. Max pooling with pool size 2 and stride 2 halves the spatial size to  $64 \times 64$ , so the output is  $64 \times 64 \times 32$ .

**Q.2** A recurrent neural network (RNN) faces the issue of vanishing gradients. Which of the following techniques addresses this problem effectively?

- A. Dropout regularization
- B. Using rectified linear units (ReLU)
- C. Gradient clipping
- D. Long Short-Term Memory (LSTM) networks

**Answer: D**

**Explanation:** LSTM networks use gated memory cells that help retain long-term information and reduce the vanishing-gradient problem in sequence learning.

**Q.3** You are training a neural network and notice that the loss decreases initially but stabilizes at a high value. What could be the most probable issue?

- A. The model is underfitting the data.
- B. The learning rate is too high.
- C. The activation function is incorrect.
- D. There is no regularization in the model.

**Answer: A**

**Explanation:** A loss that falls slightly but remains high usually means the model is not learning enough useful patterns, which is a sign of underfitting.

**Q.4** In an autonomous robot using AI and IoT, which of the following best describes how AI enhances IoT data processing?

- A. AI ensures secure data transmission over IoT networks.
- B. AI preprocesses IoT data for noise removal and missing values.
- C. AI analyzes IoT data for real-time decision-making.
- D. AI helps reduce IoT device battery consumption.

**Answer: C**

**Explanation:** AI adds intelligence to IoT systems by analyzing sensor data and supporting real-time decisions such as alerts, control, or prediction.

**Q.5** An AI-integrated IoT system detects abnormal behavior in a smart factory's machines using predictive maintenance algorithms. Which of the following models is most suitable for this task?

- A. Logistic regression
- B. K-means clustering
- C. LSTM networks
- D. Decision trees

**Answer: C**

**Explanation:** Predictive maintenance often uses time-series sensor data. LSTM networks are suitable because they can learn patterns over time.

**Q.6** Which error in IoT and robotics applications is most commonly introduced due to faulty sensor calibration?

- A. Feature scaling error
- B. Measurement noise
- C. Overfitting error
- D. Hyperparameter mismatch

**Answer: B**

**Explanation:** Faulty calibration causes sensor readings to deviate from true values, producing measurement noise or systematic measurement error.

**Q.7** In training an AI system, you observe a significant bias in predictions for a minority group. What step should you take first to address the issue?

- A. Increase the training dataset size.
- B. Balance the dataset with more samples from the minority group.
- C. Use a different activation function in the neural network.
- D. Reduce the learning rate.

**Answer: B**

**Explanation:** The first step is to improve representation by adding or balancing minority-group samples so the model learns more fairly from the data.

**Q.8** Which of the following is an example of creating an ethical AI system?

- A. Using adversarial attacks to test model robustness.
- B. Designing AI systems that prioritize transparency and interpretability.
- C. Applying aggressive data augmentation techniques.
- D. Training models on extremely large datasets.

**Answer: B**

**Explanation:** Transparency and interpretability help people understand AI decisions, which is central to responsible and ethical AI design.

**Q.9** In an object detection project using YOLOv5, you notice that small objects are often missed. Which solution would most likely improve detection performance?

- A. Increase batch size during training.
- B. Decrease the IoU threshold.
- C. Use higher-resolution images.
- D. Apply dropout regularization.

**Answer: C**

**Explanation:** Small objects occupy very few pixels. Higher-resolution images make the objects clearer and usually improve detection.

**Q.10** A chatbot struggles to respond accurately to context-based questions. What modification would most effectively address this issue?

- A. Replace the chatbot's model with a Transformer-based model.
- B. Increase the chatbot's dataset size.
- C. Add more layers to the model.
- D. Decrease the model's training epochs.

**Answer: A**

**Explanation:** Transformer-based models handle context better than simple sequence or rule-based models, making them stronger for context-dependent chatbot responses.

**Q.11** In OpenCV, which operation is used to highlight the edges of objects in an image?

- A. Gaussian blur
- B. Edge detection using the Canny method
- C. Image thresholding
- D. Morphological dilation

**Answer: B**

**Explanation:** Canny edge detection is a common OpenCV method for detecting strong edges after noise reduction and gradient analysis.

**Q.12** Which Python library is most suited for building and training convolutional neural networks?

- A. NumPy
- B. Scikit-learn
- C. TensorFlow
- D. Matplotlib

**Answer: C**

**Explanation:** TensorFlow/Keras is widely used for building and training convolutional neural networks for image tasks.

**Q.13** You are training a CNN, and during backpropagation, you notice that the weights in the first few layers are not updating. What is the most probable cause?

- A. High learning rate
- B. Exploding gradients
- C. Vanishing gradients
- D. Poor weight initialization

**Answer: C**

**Explanation:** If early layers stop receiving useful gradients, their weights may barely update. This is a classic symptom of vanishing gradients.

**Q.14** While training a neural network, the validation loss decreases initially but later starts increasing, even as training loss continues to decrease. What issue is the model facing?

- A. Underfitting
- B. Overfitting
- C. Vanishing gradients
- D. Exploding gradients

**Answer: B**

**Explanation:** Training loss falling while validation loss rises means the model is fitting the training data too closely and generalizing poorly.

**Q.15** A CNN is producing poor results on image classification tasks. Upon investigation, you find that the activation maps are too sparse. What could be the cause?

- A. Too many convolutional layers
- B. Using ReLU as the activation function
- C. Large kernel size
- D. Excessive dropout

**Answer: D**

**Explanation:** Excessive dropout can make activations too sparse and prevent the model from preserving enough useful feature information.

**Q.16** Which of the following statements about convolutional layers in CNNs is true?

- A. Convolutional layers are invariant to rotation of input features.
- B. Increasing the stride in a convolutional layer decreases the size of the feature map.
- C. Padding with "valid" mode preserves the spatial dimensions of the input.
- D. Increasing the number of convolutional filters reduces the spatial size of the feature map.

**Answer: B**

**Explanation:** Increasing stride moves the filter by larger steps, so fewer output positions are computed and the feature map becomes smaller.

**Q.17** A deep RNN is struggling with learning long-term dependencies. Which approach is most suitable to address this issue?

- A. Adding more layers to the RNN
- B. Using LSTM or GRU cells instead of standard RNN cells
- C. Applying Batch Normalization to all hidden states
- D. Reducing the size of the training dataset

**Answer: B**

**Explanation:** LSTM and GRU cells are designed to capture longer-term dependencies better than standard RNN cells.

**Q.18** Given an image classification task with an imbalanced dataset, what modification would likely result in the most robust CNN performance?

- A. Increase the kernel size of the convolutional layers
- B. Use a weighted loss function to penalize misclassifications of minority classes
- C. Add more fully connected layers to the model
- D. Use Dropout to prevent overfitting

**Answer: B**

**Explanation:** A weighted loss function gives more penalty to minority-class errors, helping the CNN learn from imbalanced data more robustly.

**Q.19** You are training a CNN and notice that the validation accuracy is consistently lower than training accuracy. Which technique is the least effective in resolving this issue?

- A. Applying data augmentation
- B. Using regularization techniques such as Dropout
- C. Adding Batch Normalization layers
- D. Increasing the number of filters in the convolutional layers

**Answer: D**

**Explanation:** Increasing filters may make the model larger and could worsen overfitting; regularization, augmentation, and normalization are more directly useful.

**Q.20** You are working on a real-world project that uses CNNs for medical image segmentation. The Dice coefficient on the validation set is low. Which approach would most likely improve performance?

- A. Use a smaller batch size during training
- B. Replace ReLU with Sigmoid as the activation function in convolutional layers
- C. Train the network with a U-Net architecture
- D. Reduce the number of pooling layers in the network

**Answer: C**

**Explanation:** U-Net is a well-known architecture for medical image segmentation because it combines localization with deep feature learning.

**Q.21** Which of the following is a primary challenge when integrating AI with IoT for smart device applications?

- A. High bandwidth requirements for edge computing
- B. Lack of sufficient IoT devices to collect data
- C. Over-reliance on supervised learning methods
- D. Difficulty in integrating IoT sensors with AI models

**Answer: D**

**Explanation:** A key practical challenge is integrating many sensor streams with AI models reliably under real-world noise, latency, and device constraints.

**Q.22** In the context of autonomous navigation for robotics, what is the role of the SLAM (Simultaneous Localization and Mapping) algorithm?

- A. Enhancing energy efficiency of IoT devices in real-time environments
- B. Enabling a robot to map an unknown environment while keeping track of its position
- C. Training deep learning models for object recognition tasks
- D. Improving communication between distributed IoT systems

**Answer: B**

**Explanation:** SLAM enables a robot to build a map of an unknown environment while simultaneously estimating its own location.

**Q.23** A robot designed for warehouse management must optimize its path while avoiding obstacles in real-time. Which AI approach is most suitable for this task?

- A. Convolutional Neural Networks (CNNs)
- B. Reinforcement Learning (RL)
- C. Generative Adversarial Networks (GANs)
- D. Unsupervised Clustering

**Answer: B**

**Explanation:** Reinforcement learning is suitable for path optimization and obstacle avoidance because an agent learns actions from rewards in an environment.

**Q.24** What is the primary advantage of implementing Federated Learning in AI-enabled IoT systems?

- A. Reduces hardware requirements for IoT devices
- B. Eliminates the need for labeled training data
- C. Preserves data privacy by training models locally on devices
- D. Improves the performance of cloud-based AI models

**Answer: C**

**Explanation:** Federated learning allows models to train locally on devices and share model updates instead of raw data, improving privacy.

**Q.25** In an AI-powered robotic arm used for assembly tasks, which sensor type is critical for detecting the force applied during operations?

- A. Infrared sensors
- B. LiDAR sensors
- C. Tactile sensors
- D. Ultrasonic sensors

**Answer: C**

**Explanation:** Tactile or force sensors are used to detect physical contact and force applied by robotic arms during manipulation.

**Q.26** Which of the following approaches is most effective in mitigating AI bias caused by underrepresented groups in training data?

- A. Using larger datasets without modifying their composition
- B. Implementing transfer learning from biased datasets
- C. Using synthetic data to supplement underrepresented groups
- D. Increasing the number of layers in the neural network

**Answer: C**

**Explanation:** Synthetic data can supplement underrepresented groups and improve model exposure to diverse examples when collected data is imbalanced.

**Q.27** Which principle is NOT part of responsible AI development?

- A. Ensuring transparency in AI decision-making processes
- B. Prioritizing the maximization of model accuracy over fairness
- C. Enforcing accountability for AI outcomes
- D. Incorporating sustainability in AI system design

**Answer: B**

**Explanation:** Responsible AI should not maximize accuracy at the expense of fairness; fairness, transparency, accountability, and sustainability matter.

**Q.28** An AI system designed for loan approval shows systematic bias against certain demographics. What is the best immediate step to address this issue?

- A. Discard the biased model and start training from scratch
- B. Add more features to the model for better accuracy
- C. Perform a fairness audit and adjust the training dataset
- D. Increase the model's complexity to better generalize

**Answer: C**

**Explanation:** A fairness audit identifies where the model is biased, and data/model adjustments can then target the actual source of unfairness.

**Q.29** How can sustainability be incorporated into the design of AI systems?

- A. Optimizing computational efficiency to reduce energy consumption
- B. Increasing model size for better accuracy
- C. Training models on cloud platforms with no energy constraints
- D. Using high-power GPUs for faster training

**Answer: A**

**Explanation:** Efficient algorithms, smaller models, and optimized computation reduce energy use and support Green AI.

**Q.30** Which of the following frameworks can help ensure ethical considerations are consistently applied in AI projects?

- A. Agile methodology
- B. ISO/IEC 27001 for Information Security
- C. IEEE Ethically Aligned Design (EAD) framework
- D. Six Sigma for quality improvement

**Answer: C**

**Explanation:** IEEE Ethically Aligned Design is specifically focused on embedding ethical considerations into autonomous and intelligent systems.

**Q.31** Which of the following techniques is best suited for improving the accuracy of an object detection model in a capstone project?

- A. Using a pre-trained object detection model like YOLO or Faster R-CNN
- B. Increasing the dataset size by adding random images
- C. Training a CNN from scratch without transfer learning
- D. Applying k-means clustering to reduce noise in the dataset

**Answer: A**

**Explanation:** Using a pre-trained detector such as YOLO or Faster R-CNN transfers learned visual features and improves object-detection accuracy with less data.

**Q.32** In an AI healthcare project for diagnosing diseases, which metric is most critical to evaluate the model's performance?

- A. Accuracy
- B. Precision
- C. Recall
- D. F1-Score

**Answer: C**

**Explanation:** In disease diagnosis, recall is critical because missing actual positive cases can be dangerous.

**Q.33** While building a chatbot for customer support, what is the primary advantage of using a transformer-based model like GPT over a rule-based chatbot?

- A. It requires no training data
- B. It can understand context and generate human-like responses
- C. It has lower computational costs
- D. It eliminates the need for a conversational dataset

**Answer: B**

**Explanation:** Transformer models can use context across a conversation and generate flexible, human-like responses, unlike rigid rule-based chatbots.

**Q.34** Which type of neural network is most appropriate for designing a real-time object detection system in a self-driving car capstone project?

- A. Recurrent Neural Networks (RNNs)
- B. Convolutional Neural Networks (CNNs)
- C. Generative Adversarial Networks (GANs)
- D. Transformer Networks

**Answer: B**

**Explanation:** CNNs are the core neural network type for extracting visual features used in real-time object detection systems.

**Q.35** In an AI capstone project involving sentiment analysis for chatbot improvement, which preprocessing step is most critical for handling input text?

- A. Normalizing pixel values
- B. Tokenizing and removing stopwords
- C. Applying Fourier transforms
- D. Scaling numerical values to a specific range

**Answer: B**

**Explanation:** For sentiment analysis, text must first be split into tokens and cleaned so that words can be analyzed by the model.

**Q.36** In OpenCV, which function is most commonly used for edge detection, and what is its key advantage?

- A. `cv2.HoughLines()`; it accurately detects circular edges
- B. `cv2.Canny()`; it reduces noise and detects strong edges
- C. `cv2.Laplacian()`; it computes gradients across all directions
- D. `cv2.findContours()`; it detects object boundaries efficiently

**Answer: B**

**Explanation:** `cv2.Canny()` is commonly used for edge detection; it reduces noise and identifies strong edges through a multi-stage process.

**Q.37** When designing a CNN for image classification, what is the purpose of adding pooling layers, such as max pooling?

- A. To increase the number of parameters in the network
- B. To reduce the spatial dimensions of feature maps
- C. To prevent overfitting during training
- D. To convert the input data into grayscale

**Answer: B**

**Explanation:** Pooling layers reduce the height and width of feature maps, lowering computation while retaining important features.

**Q.38** Which Python library is most suitable for performing data augmentation in deep learning projects, and what is one of its common techniques?

- A. NumPy; element-wise operations
- B. Matplotlib; plotting transformed data
- C. TensorFlow/Keras; image flipping and rotation
- D. Pandas; tabular data reshaping

**Answer: C**

**Explanation:** TensorFlow/Keras provides image preprocessing and augmentation utilities such as flipping, rotation, shifting, and zooming.

**Q.39** In OpenCV, what is the output of the `cv2.findContours()` function, and how can it be used in object detection?

- A. A list of contour points; it helps identify object shapes in an image
- B. A heatmap of detected objects; it highlights regions of interest
- C. A set of bounding boxes; it directly detects object locations
- D. A binary mask of the image; it filters out noise

**Answer: A**

**Explanation:** `findContours` returns contour point lists that describe object boundaries and can be used to identify shapes.

**Q.40** When implementing a CNN using Keras, which activation function is most commonly used in hidden layers, and why?

- A. Sigmoid; it prevents vanishing gradients
- B. Tanh; it ensures values remain centered around zero
- C. ReLU; it introduces non-linearity and avoids vanishing gradients
- D. Softmax; it normalizes output for classification

**Answer: C**

**Explanation:** ReLU is commonly used in hidden layers because it introduces non-linearity and reduces vanishing-gradient issues compared with saturated activations.

## Section B: Achievers Section

**Q.41** Which of the following decorators in Python is used to define a class method?

- A. @staticmethod
- B. @classmethod
- C. @property
- D. @abstractmethod

**Answer: B**

**Explanation:** @classmethod defines a method that receives the class as the first argument rather than an object instance.

**Q.42** What will be the output of the following code snippet?

```
def func(a, b=[]):
```

```
    b.append(a)
```

```
    return b
```

```
    print(func(1))
```

```
    print(func(2))
```

- A. [1], [2]
- B. [1], [1, 2]
- C. [1, 2], [1, 2]
- D. Raises an error

**Answer: B**

**Explanation:** The default list argument is shared across calls, so the second call appends 2 to the same list that already contains 1.

**Q.43** In Python's multiprocessing module, which method is used to share data between processes?

- A. Queue
- B. Lock
- C. Array
- D. Manager

**Answer: D**

**Explanation:** A multiprocessing Manager can create shared objects such as shared lists or dictionaries that multiple processes can access.

**Q.44** What does the following generator expression produce?

```
gen = (x**2 for x in range(3))
print(next(gen))
print(next(gen))
print(next(gen))
print(next(gen))
```

- A. 0, 1, 4, StopIteration
- B. 1, 4, StopIteration
- C. Raises TypeError
- D. 0, 1, 4, None

**Answer: A**

**Explanation:** The generator produces 0, then 1, then 4. A fourth next() call raises StopIteration because the generator is exhausted.

**Q.45** What is the purpose of the zip\_longest() function in the itertools module?

- A. To combine two iterables, truncating to the length of the shorter iterable
- B. To combine two iterables, filling missing values with a specified value
- C. To create combinations of elements from two iterables
- D. To merge two sorted iterables

**Answer: B**

**Explanation:** zip\_longest combines iterables and fills missing values when one iterable is shorter than the other.

**Q.46** Which Python standard library module can be used to work with low-level file descriptors?

- A. os
- B. io
- C. sys
- D. fileinput

**Answer: A**

**Explanation:** The os module provides functions for low-level operating-system interfaces including file descriptors.

**Q.47** What is the primary difference between is and == in Python?

- A. is checks for equality, while == checks for object identity
- B. is checks for object identity, while == checks for equality
- C. Both perform the same operation
- D. is is used only for strings

**Answer: B**

**Explanation:** is checks whether two references point to the same object; == checks whether their values are equal.

**Q.48** What does the following Python code do?

```
import collections
Counter = collections.Counter
data = [1, 2, 2, 3, 3, 3]
result = Counter(data).most_common(2)
print(result)
```

- A. Prints the two most common elements in the list
- B. Sorts the list in descending order
- C. Removes duplicates from the list
- D. Raises a TypeError

**Answer: A**

**Explanation:** Counter(data).most\_common(2) returns the two most frequent values and their counts.

**Q.49** What will be the output of the following Python code?

```
class A:
    def __init__(self):
        self.x = 1
class B(A):
    def __init__(self):
        super().__init__()
        self.y = 2
        b = B()
        print(b.x, b.y)
```

A. Raises an AttributeError  
 B. 1, 2  
 C. None, 2  
 D. 0, 2

**Answer: B**

**Explanation:** B calls the parent constructor through `super().__init__()`, setting `x = 1`, and then sets `y = 2`.

**Q.50** Which of the following Python constructs is used to handle errors during code execution?

- A. assert
- B. try...except
- C. raise
- D. del

**Answer: B**

**Explanation:** try...except catches and handles errors raised during program execution.

## Consolidated Answer Key

Q	Ans	Q	Ans	Q	Ans	Q	Ans	Q	Ans
1	A	2	D	3	A	4	C	5	C
6	B	7	B	8	B	9	C	10	A
11	B	12	C	13	C	14	B	15	D
16	B	17	B	18	B	19	D	20	C
21	D	22	B	23	B	24	C	25	C
26	C	27	B	28	C	29	A	30	C
31	A	32	C	33	B	34	B	35	B
36	B	37	B	38	C	39	A	40	C
41	B	42	B	43	D	44	A	45	B
46	A	47	B	48	A	49	B	50	B

