

SCO INTERNATIONAL AI OLYMPIAD

CLASS 12 QUESTION PAPER - SET S

2025-26 | Answer Key and Detailed Explanations Included

Designed for high-school learners preparing for AI, machine learning, computer vision, robotics, IoT, and responsible AI reasoning.

- professionally formatted question blocks with code snippets, answer keys, and explanations
- scenario-based AI ethics, sustainability, robotics, OpenCV, Python, and deep learning practice
- publication-ready layout for website PDF viewing and downloadable learning support

Deep Learning

Robotics

IoT

NLP

Ethics

OpenCV

Python

Data Science

Sustainability

ML Practice

Candidate Guidelines

Class 12	Total Questions 50	Time 1 Hour
Name:		Registration ID:
School / Institution:		Contact No.:

1. Before the exam begins, candidates may receive additional time to complete the OMR sheet or student information section as instructed by the invigilator.
2. Write the name, registration ID, school or institution details, class, roll number, and contact number clearly wherever applicable.
3. There are 50 questions. Each question has only one correct answer. Select one response choice only.
4. Calculator use is not allowed unless separately announced by the official exam instructions.
5. Use an HB pencil or a blue/black ballpoint pen to mark answers on the OMR sheet or as directed by the exam platform.
6. For the sample item: What is the capital of India? A. Kolkata B. Mumbai C. Delhi D. Chennai. The correct response is C.
7. Submit the answer sheet or complete the online submission only after reviewing all responses.

Deep Learning Basics

Focus: neural network foundations, CNN/RNN concepts, feature extraction, preprocessing, and model purpose.

Q1

What is the fundamental building block of a deep neural network?

A	A database
B	A neuron
C	A decision tree
D	A support vector machine

Answer: B) A neuron

Explanation: In deep learning, the neuron (or node) is the basic unit that processes inputs by applying weights, biases, and an activation function to produce an output.

Q2

Which process adjusts the weights in a neural network based on the error of its predictions?

A	Forward propagation
B	Activation
C	Backpropagation
D	Data normalization

Answer: C) Backpropagation

Explanation: Backpropagation is the algorithm that computes the gradient of the loss function with respect to each weight, allowing the network to update its parameters to reduce errors.

Q3

Consider the following Python code snippet:


```
import numpy as np

inputs = np.array([1.0, 2.0, 3.0])
weights = np.array([0.5, -0.6, 0.1])
bias = 0.2

output = np.dot(inputs, weights) + bias
print(output)
```

What is the output printed by this code?

Weighted-sum neuron model



A	0.2
B	-0.2
C	1.0
D	2.0

Answer: B) -0.2

Explanation: The dot product is $(1.0 \times 0.5) + (2.0 \times -0.6) + (3.0 \times 0.1) = 0.5 - 1.2 + 0.3 = -0.4$. Adding the bias gives $-0.4 + 0.2 = -0.2$. Therefore, the printed output is -0.2.

Q4

In Convolutional Neural Networks (CNNs), which layer is typically used to reduce the spatial dimensions of an image?

A	Convolutional layer
B	Pooling layer

C	Fully connected layer
D	Recurrent layer

Answer: B) Pooling layer

Explanation: Pooling layers (e.g., max pooling or average pooling) reduce the spatial dimensions (width and height) of the input, helping to decrease the number of parameters and computation in the network.

Q5

What is the primary purpose of a convolutional filter in a CNN?

A	To reduce overfitting
B	To extract features from the input image
C	To convert text into numerical data
D	To perform classification directly

Answer: B) To extract features from the input image

Explanation: Convolutional filters slide over the input image to detect features such as edges, textures, and patterns, which are essential for tasks like image classification.

Q6

Recurrent Neural Networks (RNNs) are best suited for processing which type of data?

A	Static images
B	Sequential data
C	Unstructured text without context
D	Isolated numerical values

Answer: B) Sequential data

Explanation: RNNs are designed to handle sequential data (e.g., time series, sentences, speech) because they maintain a memory of previous inputs, enabling them to capture temporal dependencies.

Q7

Consider the following TensorFlow/Keras code snippet:

```
import tensorflow as tf
from tensorflow.keras.models import Sequential
from tensorflow.keras.layers import Dense

model = Sequential()
model.add(Dense(5, activation='relu', input_shape=(3,)))
model.add(Dense(1, activation='sigmoid'))

model.compile(optimizer='adam', loss='binary_crossentropy')
print(model.summary())
```

What type of problem is this neural network most likely designed for?

A	Multi-class classification
B	Binary classification
C	Regression
D	Clustering

Answer: B) Binary classification

Explanation: The network uses a sigmoid activation function in the output layer and binary crossentropy as the loss function, which are typically used for binary classification tasks (i.e., deciding between two classes).

Q8

What does the term “data preprocessing” refer to in the context of AI and machine learning?

A	Training the model on raw data
B	Cleaning and transforming raw data into a format suitable for modeling

C	Visualizing data using graphs and charts
D	Deploying the model into production

Answer: B) Cleaning and transforming raw data into a format suitable for modeling

Explanation: Data preprocessing involves tasks such as cleaning, normalization, transformation, and feature extraction, all of which prepare the raw data for effective model training.

Q9

In computer vision applications, what is the primary goal of using a CNN?

A	To generate natural language from images
B	To classify images or detect objects within images
C	To process sequential video data
D	To compress image files

Answer: B) To classify images or detect objects within images

Explanation: CNNs are highly effective at learning spatial hierarchies and patterns in image data, making them ideal for image classification and object detection tasks.

Q10

Which of the following is a major advantage of using deep learning for image processing?

A	It requires no data preprocessing
B	It automatically extracts relevant features from images without manual intervention
C	It only works with grayscale images

D

It eliminates the need for any human supervision during training

Answer: B) It automatically extracts relevant features from images without manual intervention**Explanation:** One of the key benefits of deep learning, particularly CNNs, is their ability to learn and extract hierarchical features directly from raw image data, reducing the need for manual feature engineering.

AI in Robotics, IoT, and NLP Systems

Focus: edge computing, sensor fusion, smart systems, chatbots, reinforcement learning, and autonomous robotics.

Q11

In a smart manufacturing system, IoT sensors collect real-time data (e.g., temperature, vibration) from machinery. Which of the following best describes how AI is used to improve maintenance in such a system?

A

AI simply monitors sensor data and triggers an alarm when a fixed threshold is exceeded.

B

AI analyzes both historical and real-time sensor data to predict potential equipment failures before they occur.

C

AI replaces the need for sensors by simulating machine behavior based on static rules.

D


AI directly controls the mechanical components without any data analysis.

Answer: B**Explanation:** In smart manufacturing, AI algorithms (often using machine learning) analyze historical and real-time sensor data to identify patterns indicative of impending equipment failures. This predictive maintenance approach minimizes downtime and reduces costs. Option A oversimplifies the process; options C and D ignore the critical role of continuous data analysis.

Q12

Edge computing is increasingly used in IoT deployments to reduce latency in decision-making. Which statement best explains the benefit of edge computing when integrated with AI for robotics applications?

AI + IoT + Robotics: low-latency flow



Processing near the data source reduces latency.

A	It centralizes data in remote data centers to leverage high processing power.
B	It processes data near the source, thereby reducing latency and enabling real-time responses.
C	It focuses solely on encrypting data for improved security in IoT networks.
D	It replaces AI algorithms with simpler heuristics to speed up processing.

Answer: B

Explanation: Edge computing brings data processing closer to the data source (e.g., on the robot or a nearby gateway), reducing transmission delays and enabling quick, real-time decisions. This is critical for robotics applications where response time is essential. Option A describes cloud computing, option C is a partial benefit, and option D misrepresents edge computing’s purpose.

Q13

Sensor fusion is a vital technique in AI-driven robotics. What is the primary purpose of sensor fusion in these systems?

A	To eliminate redundant sensor data and rely on a single sensor.
B	To combine data from multiple sensors, enhancing the robot’s perception of its environment.

C	To display sensor data visually for human operators without affecting robot decision-making.
D	To reduce hardware costs by removing the need for various sensors.

Answer: B

Explanation: Sensor fusion integrates data from multiple sensors (e.g., cameras, LIDAR, ultrasonic sensors) to provide a more accurate and robust understanding of the environment. This enhanced perception is essential for tasks such as navigation and object detection. Option A discards useful data, while C and D do not capture the central role of sensor fusion in improving decision-making.

Q14

Modern chatbots rely on several natural language processing (NLP) components. Which component is primarily responsible for determining the purpose behind a user's message?

A	Entity Extraction
B	Intent Recognition
C	Sentiment Analysis
D	Response Generation

Answer: B

Explanation: Intent recognition is the process by which a chatbot interprets the underlying purpose or goal of a user's input. While entity extraction (option A) identifies key elements and sentiment analysis (option C) gauges the emotion, intent recognition drives the decision-making process on how to respond. Response generation (option D) occurs after intent is understood.

Q15

An organization deploys an AI-powered chatbot across several messaging platforms (e.g., Facebook Messenger, WhatsApp, and Slack). What is a major challenge in ensuring a consistent user experience?

A	Ensuring uniform response time across platforms.
---	--

B	Training the chatbot to support multiple languages.
C	Integrating the chatbot with diverse platform-specific features and user interfaces.
D	Storing conversation logs in a single centralized database.

Answer: C

Explanation: Different messaging platforms have unique constraints, features, and user interface designs. Ensuring the chatbot adapts appropriately to each platform is a significant challenge, as it must maintain consistency in functionality and user experience. While A and B can be concerns, option C directly addresses the complexity of platform integration.

Q16

Reinforcement learning (RL) is frequently applied to improve the behavior of autonomous robots over time. Which statement best describes the advantage of using RL in robotics?

A	RL allows robots to imitate human behavior through direct instruction.
B	RL enables robots to learn optimal actions by receiving rewards or penalties based on their actions.
C	RL uses fixed rules to control robot actions in all scenarios.
D	RL eliminates the need for sensor data by relying entirely on pre-programmed responses.

Answer: B

Explanation: Reinforcement learning is a technique where robots (agents) learn from interactions with the environment by receiving feedback (rewards or penalties). This trial-and-error approach allows them to optimize actions for complex tasks. Option A describes imitation learning, while options C and D contradict the adaptive and data-driven nature of RL.

Q17

What is a “digital twin” in the context of integrating AI with IoT and robotics?

A	A virtual replica of a physical system that enables simulation, prediction, and optimization of real-world performance.
B	Two identical robots operating together to double system efficiency.
C	A backup IoT device that activates in case the primary device fails.
D	An AI algorithm that clones human decision-making in robotic systems.

Answer: A

Explanation: A digital twin is a virtual model of a physical object or system that receives real-time data from IoT sensors. AI algorithms use this model to simulate behavior, predict outcomes, and optimize performance. Option B misunderstands the concept, while options C and D do not accurately reflect the digital twin’s purpose.

Q18

When integrating AI with IoT devices in a smart home environment, why is data preprocessing a critical step?

A	It encrypts all incoming data to ensure privacy.
B	It transforms raw sensor data into a clean, structured format suitable for efficient analysis by AI algorithms.
C	It archives raw data indefinitely for potential future use.
D	It directly generates control commands for IoT devices.

Answer: B

Explanation: Data preprocessing involves cleaning, normalizing, and transforming raw data into a format that can be efficiently and accurately processed by AI algorithms. This step is essential to improve model performance and reliability. Option A is related to security, option C is about data storage, and option D mischaracterizes the purpose of preprocessing.

Q19

Which type of neural network architecture is best suited for handling sequential language data in AI-powered chatbot systems?

A	Convolutional Neural Networks (CNNs)
B	Recurrent Neural Networks (RNNs)
C	Generative Adversarial Networks (GANs)
D	Feedforward Neural Networks

Answer: B

Explanation: Recurrent Neural Networks (RNNs), including variants like LSTM (Long Short-Term Memory) networks and GRUs (Gated Recurrent Units), are designed to handle sequential data by maintaining internal states. This makes them particularly effective for natural language processing tasks in chatbots. CNNs are more commonly used for image processing, GANs for data generation, and feedforward networks lack temporal context.

Q20

Autonomous robots often rely on computer vision for navigation in dynamic environments. Which combination of AI techniques is most effective for enabling a robot to identify obstacles and plan a safe path?

A	Rule-based algorithms combined with static image recognition.
B	Deep learning for object detection (using CNNs) and reinforcement learning for adaptive decision-making.
C	Traditional machine learning with manual sensor calibration.
D	Pre-programmed paths without real-time adaptation.

Answer: B

Explanation: Deep learning methods, such as Convolutional Neural Networks (CNNs), enable robust object detection and scene understanding. When combined with reinforcement learning, which allows the robot to learn from its environment and adapt its

navigation strategy, the robot can effectively handle dynamic, unpredictable environments. Option A is too rigid, option C does not fully leverage the potential of modern AI, and option D lacks the necessary adaptability.

Responsible AI, Sustainability, and Ethics

Focus: data bias, fairness, privacy, consent, sustainable AI practices, and real-world governance dilemmas.


Q21

In the city of Innovata, the police department recently deployed an AI-based facial recognition system. During testing, they discovered that the system was less accurate in identifying individuals with darker skin tones. The team realized that the training data had mostly light-skinned faces, leading to unintended bias.

Question:

Which of the following is the most likely reason for the system’s poor performance on darker skin tones?

Responsible AI: representative data matters



A	The algorithm’s architecture is inherently biased.
B	There was insufficient diverse training data.
C	The system was programmed to ignore dark features.
D	The camera quality was too low.

Answer: B

Explanation: The system’s bias is most likely due to an imbalanced training dataset that did not include enough diverse examples. Algorithms learn from the data provided; if that data lacks diversity, the resulting model will also be biased. Option A is incorrect because the architecture is not inherently biased, and options C and D do not address the data imbalance issue.

Q22

GreenTech Solutions developed an AI-powered sensor network to monitor air quality and energy usage in urban areas. A Python script aggregates real-time sensor data and forecasts pollution levels. This system helps city planners implement timely interventions to reduce environmental impact and save energy.

```
import numpy as np

# Simulated sensor data: [CO2, NO2, O3]
sensor_data = np.array([400, 30, 50])
weights = np.array([0.6, 0.3, 0.1])
pollution_score = np.dot(sensor_data, weights)
print("Pollution Score:", pollution_score)
```

Question:

What is one sustainability benefit of using AI in this environmental monitoring system?

A	It eliminates the need for physical sensors.
B	It enables proactive measures by forecasting pollution levels.
C	It reduces the cost of data storage.
D	It automatically repairs faulty sensors.

Answer: B

Explanation: The AI system forecasts pollution levels based on sensor data, allowing city planners to take proactive measures (such as traffic regulation or industrial adjustments) to improve air quality and sustainability. Options A, C, and D do not accurately describe the primary benefit of this AI application.

Q23

A large corporation uses an AI system to screen job applications. Soon, HR notices that candidates from certain backgrounds are consistently ranked lower. On investigation, it is discovered that historical hiring data—which contained biased human decisions—was used to train the system.

Question:

What ethical issue does this scenario highlight?

A	Data privacy breach
B	Algorithmic bias due to historical data
C	Inaccurate job role matching
D	Over-reliance on technical interviews

Answer: B

Explanation: The scenario clearly demonstrates algorithmic bias: the AI learned from historical data that contained human biases. This results in the system unfairly penalizing candidates from certain backgrounds. Options A, C, and D are not the central ethical concerns in this context.

Q24

A university research team is training a deep neural network for image recognition. They notice that each training session consumes a significant amount of energy, leading to high costs and a larger carbon footprint. The team debates ways to optimize their models and training processes to be more environmentally friendly.

Question:

Which strategy is most likely to improve sustainability during model training?

A	Increasing the model size to achieve better accuracy.
B	Reducing training epochs and using transfer learning.
C	Ignoring model optimization to save time.
D	Training multiple models in parallel without efficiency concerns.

Answer: B

Explanation: Using transfer learning (leveraging pre-trained models) and reducing the number of training epochs can significantly decrease energy consumption without sacrificing accuracy. This strategy helps lower the carbon footprint associated with extensive training. Option A would likely worsen energy consumption, and C and D do not address sustainability concerns.

Q25

A multinational bank deploys an AI-powered chatbot to assist customers. However, some users report that the chatbot provides less helpful responses to non-native English speakers. The development team discovers that the training dataset predominantly consists of native English language queries, resulting in less effective natural language processing for diverse accents and dialects.

Question:

What is the primary ethical concern in this scenario?

A	Privacy violation
B	Lack of personalization

C	Unfair treatment due to biased training data
D	Slow response time

Answer: C

Explanation: The main issue is the biased training data, which causes the chatbot to perform less effectively for non-native English speakers. This results in unfair treatment and a suboptimal customer experience. Options A, B, and D do not capture the core ethical concern of bias in data representation.

Q26

A fintech company uses an AI system to automate loan approvals. After a period of operation, it is observed that applicants from certain zip codes are being denied loans at a higher rate. An investigation reveals that the training data included historical loan decisions influenced by socio-economic biases.

Question:

Which measure would most likely improve fairness in this loan approval system?

A	Removing zip code data from the input features
B	Increasing the complexity of the model
C	Reducing the number of applicants
D	Deploying the system on more powerful servers

Answer: A

Explanation: Removing or carefully treating zip code data can reduce reliance on a feature that may act as a proxy for socio-economic bias. In a real system, this should be combined with fairness audits, representative data, proxy-feature review, and human oversight.

Q27

A research group is analyzing a dataset used to train an AI system for college admissions. They suspect the dataset might be biased against certain groups. They write a Python script to compute the distribution of acceptance rates by demographic groups.

```
import pandas as pd
```

```
# Sample data: each row represents an applicant with their demographic group and admission status (1
for admitted, 0 for not admitted)
data = pd.DataFrame({
    'group': ['A', 'B', 'A', 'B', 'A', 'B', 'A', 'B'],
    'admitted': [1, 0, 1, 0, 0, 0, 1, 1]
})

acceptance_rates = data.groupby('group')['admitted'].mean()
print(acceptance_rates)
```

Question:

If group A has an acceptance rate of 75% and group B has an acceptance rate of 25%, what is the primary ethical issue highlighted by the script's output?

A	Data overfitting
B	Algorithmic bias
C	Inaccurate data collection
D	Model underfitting

Answer: B

Explanation: The script reveals a significant disparity in acceptance rates between groups, indicating potential algorithmic bias. The AI system might be favoring one group over another due to biased training data or features. Options A, C, and D do not capture the fairness and bias concerns raised by the data analysis.

Q28

EcoGrid Inc. uses AI to optimize energy distribution in a smart grid. The system predicts peak usage times and adjusts the distribution accordingly, reducing energy waste. This not only lowers costs but also minimizes the environmental impact by decreasing unnecessary power generation.

Question:

What is the main sustainability advantage of using AI in this smart grid application?

A	It increases the complexity of power systems.
B	It enables efficient energy distribution and reduces waste.
C	It requires more energy for data processing.
D	It relies on outdated forecasting methods.

Answer: B

Explanation: The AI system’s ability to predict peak usage and optimize energy distribution helps reduce waste and improves the overall efficiency of the smart grid, leading to lower environmental impact and enhanced sustainability. Options A, C, and D contradict the benefits described in the scenario.

Q29

A school implements an AI system to monitor student attendance and behavior via classroom cameras and sensors. While the system helps identify safety issues, parents raise concerns about the invasion of privacy and the continuous surveillance of their children.

Question:

Which ethical principle is most at risk in this scenario?

A	Fairness
B	Transparency
C	Privacy
D	Efficiency

Answer: C

Explanation: The scenario raises significant privacy concerns due to the continuous monitoring and data collection of students. Balancing safety with individual privacy rights is a major ethical challenge in deploying such AI systems. Options A, B, and D are not the central issue in this case.

Q30

A tech company developing autonomous vehicles faces a moral dilemma. In an unavoidable accident scenario, the car’s AI must choose between two actions: one that minimizes harm to pedestrians but may injure the passengers, or one that protects the passengers but increases risk to pedestrians. This ethical dilemma is known as the “trolley problem” for autonomous vehicles.

Question:

What is the key ethical challenge presented in this scenario?

A	Ensuring high-speed performance
B	Balancing passenger safety against pedestrian safety
C	Optimizing fuel efficiency

D

Improving sensor accuracy

Answer: B

Explanation: The scenario encapsulates the classic ethical dilemma of balancing the safety of passengers versus pedestrians. Autonomous vehicle decision-making in such critical situations requires ethical frameworks that weigh different risks, making it a complex challenge. Options A, C, and D do not address the ethical conflict described.

Computer Vision and OpenCV Applications

Focus: image reading, color conversion, resizing, edge detection, face/object detection, webcam frames, OCR preparation, and privacy controls.

Q31

In a capstone project, a team is using OpenCV to build an object detection system. The first step is reading an image from disk. Consider the following code snippet:

```
import cv2
image = cv2.imread('scene.jpg')
if image is None:
    print("Error loading image!")
```

else:

```
    print("Image loaded successfully.")
```

Question:

Which of the following statements about the function used is correct?

A	cv2.imread() reads an image in grayscale by default.
B	cv2.imread() returns None if the file cannot be found.
C	cv2.imread() automatically resizes the image to 256×256 pixels.
D	cv2.imread() converts the image to RGB format.

Answer: B

Explanation: The function cv2.imread() returns None if the specified file cannot be found or loaded. By default, it reads images in color (BGR order) and does not resize them automatically. Option A is false (default is color), option C is incorrect (no automatic resizing), and option D is incorrect (OpenCV uses BGR, not RGB).


Q32

An AI project involves analyzing road signs from images using edge detection. The team uses the following code snippet for detecting edges:

```
import cv2
gray = cv2.cvtColor(image, cv2.COLOR_BGR2GRAY)
edges = cv2.Canny(gray, 100, 200)
cv2.imshow('Edges', edges)
cv2.waitKey(0)
cv2.destroyAllWindows()
```

Question:

What is the primary purpose of the cv2.Canny() function in this context?

Computer vision preprocessing pipeline

Clean preprocessing improves reliable inference.

A	To blur the image for noise reduction.
B	To detect the boundaries of objects by finding intensity gradients.
C	To convert the image to binary format.
D	To equalize the histogram of the image.

Answer: B

Explanation: The cv2.Canny() function is used to perform Canny edge detection, which identifies edges by detecting regions with strong intensity gradients. It does not blur, binarize, or equalize histograms; those are separate processing steps.

Q33

In a real-life AI project, a student must convert an image's color space for proper visualization in a web application. The code snippet is as follows:

```
import cv2
image = cv2.imread('scene.jpg')
converted = cv2.cvtColor(image, cv2.COLOR_BGR2RGB)
```

Question:

Why is the color space conversion from BGR to RGB necessary in many AI applications?

A	Because BGR images have higher resolution than RGB images.
B	Because many libraries (e.g., Matplotlib) assume images are in RGB format for correct display.
C	Because converting to RGB reduces the image file size.
D	Because BGR cannot be processed by neural networks.

Answer: B

Explanation: OpenCV loads images in BGR format by default, while many visualization libraries (like Matplotlib) expect images in RGB format. Converting ensures correct color display. Options A, C, and D are not true.

Q34

For an AI project that involves training a deep learning model on image data, a student needs to resize images to a uniform size. Consider the following code snippet:

```
import cv2
resized = cv2.resize(image, (224, 224), interpolation=cv2.INTER_AREA)
```

Question:

Which statement best describes the role of the interpolation parameter used above?

A	It defines the color space for the resized image.
B	It determines the method used to calculate pixel values when resizing.
C	It automatically increases image brightness during resizing.
D	It converts the image from 2D to 3D.

Answer: B

Explanation: The interpolation parameter in `cv2.resize()` specifies the algorithm used to interpolate pixel values when resizing an image. `cv2.INTER_AREA` is generally used for image shrinking, offering good results. The other options are not relevant to interpolation.

Q35

A capstone project involves building a face detection system using OpenCV's Haar cascade classifiers. During the training phase, the team needs to collect and label a large dataset of face and non-face images.

Question:

Which of the following is the most critical step in ensuring that the Haar cascade classifier is robust?

A	Using only a single image per person to train the model.
B	Balancing the dataset with a diverse set of positive (face) and negative (non-face) images.
C	Training exclusively on synthetic images generated by software.
D	Ignoring variations in lighting and background in the training images.

Answer: B

Explanation: A robust Haar cascade classifier requires a diverse and balanced dataset of positive and negative examples to learn the features of faces under various conditions. Using only one image or synthetic images, or ignoring real-world variations, will lead to poor generalization.

Q36

In a real-life project, an engineering team integrates a pre-trained deep learning model for object detection with OpenCV. They load a model using the following code snippet:

```
import cv2
net = cv2.dnn.readNetFromCaffe('deploy.prototxt', 'model.caffemodel')
blob = cv2.dnn.blobFromImage(image, scalefactor=0.007843, size=(300, 300), mean=(127.5, 127.5, 127.5))
net.setInput(blob)
detections = net.forward()
```

Question:

What is the purpose of the `cv2.dnn.blobFromImage()` function in this context?

A	To directly display the image on screen.
B	To convert the image into a format suitable for processing by the deep learning model.

C	To perform image thresholding.
D	To save the image to disk.

Answer: B

Explanation: cv2.dnn.blobFromImage() preprocesses the image (scaling, resizing, mean subtraction) to create a “blob” that can be used as input to a deep learning network. This function standardizes the image for the model's expected input format.

Q37

A team is developing an interactive installation that uses a webcam to track objects in real time. They write the following code to capture video frames:

```
import cv2
cap = cv2.VideoCapture(0)
while True:
    ret, frame = cap.read()
    if not ret:
        break
    cv2.imshow('Frame', frame)
    if cv2.waitKey(1) & 0xFF == ord('q'):
        break
```

```
cap.release()
```

```
cv2.destroyAllWindows()
```

Question:

What is the role of cv2.waitKey(1) in the video capture loop?

A	It delays the loop for 1 second between frames.
B	It captures a key event for 1 millisecond, allowing for real-time display updates and break condition checking.
C	It saves each frame to disk every 1 millisecond.
D	It converts the frame to grayscale.

Answer: B

Explanation: cv2.waitKey(1) waits for a key event for 1 millisecond. This short delay allows the frame to be displayed in real time and checks if the 'q' key has been pressed to exit the loop. It does not delay for a full second, save frames, or convert the frame's color space.

Q38

A capstone project focuses on license plate recognition using OpenCV and machine learning. After detecting the license plate region, the team applies thresholding and morphological operations to extract individual characters for optical character recognition (OCR).

Question:

Which of the following steps is most critical for preparing the extracted license plate region for OCR?

A	Applying a Gaussian blur to enhance edges.
B	Converting the region to grayscale and applying thresholding to create a binary image.
C	Increasing the contrast by a fixed percentage.
D	Rotating the image by 90 degrees.

Answer: B

Explanation: Converting the license plate region to grayscale and applying thresholding to generate a binary image is crucial for OCR, as it simplifies the image and highlights character boundaries. Options A, C, and D are less directly beneficial for OCR performance.

Q39

In a real-life AI project, a team uses OpenCV to preprocess images before feeding them into a TensorFlow-based classifier. They utilize the following code:

```
import cv2
import numpy as np
import tensorflow as tf

image = cv2.imread('object.jpg')
resized = cv2.resize(image, (224, 224))
normalized = resized / 255.0
input_tensor = np.expand_dims(normalized, axis=0)
predictions = tf.keras.models.load_model('model.h5').predict(input_tensor)
print("Predictions:", predictions)
```

Question:

What is the main purpose of normalizing the image by dividing pixel values by 255.0?

A	To convert the image from BGR to RGB.
B	To scale pixel values to a [0,1] range, which is commonly expected by deep learning models.

C	To reduce the image dimensions.
D	To apply a blur effect.

Answer: B

Explanation: Dividing pixel values by 255.0 scales them from their original range of 0–255 to a normalized range of 0–1. This normalization is a standard preprocessing step for many deep learning models, ensuring numerical stability during training and inference. Options A, C, and D are not achieved by this operation.

Q40

A tech startup is building an interactive kiosk that uses facial recognition via OpenCV to personalize customer experiences. While the technology works well, the team must also consider ethical implications such as data privacy and consent.

Question:

Which of the following measures is most important to address ethical concerns in this scenario?

A	Increasing the resolution of the cameras for better recognition accuracy.
B	Implementing data anonymization and obtaining explicit consent from users before collecting facial data.
C	Reducing the number of kiosks to minimize data collection.
D	Sharing all collected facial data with third parties to improve system performance.

Answer: B

Explanation: Implementing data anonymization and obtaining explicit consent are essential steps in protecting user privacy and ensuring ethical use of facial recognition technology. This addresses the core ethical concern of data privacy. Options A and C do not directly mitigate privacy risks, while option D violates ethical guidelines.

Python, Data Science, and Machine Learning Practice

Focus: Python logic, Pandas/NumPy outputs, CNN layer behavior, decorators, regression, video frames, and data augmentation.

Q41

In a capstone project, a data scientist is analyzing sales data. They use a Pandas DataFrame to group sales by product category and then compute the total revenue for each group. Consider the following code:

```
import pandas as pd

data = {
    'category': ['electronics', 'clothing', 'electronics', 'furniture', 'clothing'],
    'revenue': [200, 150, 300, 400, 250]
}

df = pd.DataFrame(data)
result = df.groupby('category')['revenue'].apply(lambda x: x.sum())
print(result)
```

Question:

What is the output of this code?

A	<pre>category clothing 400 electronics 500 furniture 400 Name: revenue, dtype: int64</pre>
B	<pre>category clothing 250 electronics 300 furniture 400 Name: revenue, dtype: int64</pre>
C	<pre>clothing 400 electronics 500 furniture 400 dtype: int64</pre>
D	An error occurs due to incorrect lambda usage.

Answer: A

Explanation: For clothing: $150 + 250 = 400$; electronics: $200 + 300 = 500$; furniture: 400. Option A correctly shows the grouping by category with the sum of revenues. The lambda function is valid, so no error occurs.

Q42

A developer is processing a list of numbers and wants to square only the even numbers while keeping odd numbers unchanged. They write the following code:

```
numbers = [1, 2, 3, 4, 5, 6]
result = [x**2 if x % 2 == 0 else x for x in numbers]
print(result)
```

Question:

What is the printed output?

A	[1, 4, 3, 16, 5, 36]
B	[1, 2, 3, 4, 5, 6]
C	[1, 4, 9, 16, 25, 36]
D	[1, 4, 3, 8, 5, 12]

Answer: A

Explanation: Even numbers 2, 4, and 6 are squared ($2^2=4$, $4^2=16$, $6^2=36$) while odd numbers remain unchanged. Thus, the list becomes [1, 4, 3, 16, 5, 36].

Q43

For an image processing project, a student needs to load an image and convert it to grayscale for further analysis. They use the following code:

```
import cv2
image = cv2.imread('color_image.jpg')
gray = cv2.cvtColor(image, cv2.COLOR_BGR2GRAY)
print(gray.shape)
```

Question:

Assuming the original image has dimensions 480×640 with 3 channels, what does the output represent?

A	(480, 640, 3)
B	(480, 640)
C	(640, 480, 1)
D	(640, 480)

Answer: B**Explanation:** After converting to grayscale, the image loses its channel dimension. Thus, if the original dimensions are (480, 640, 3), the grayscale image shape becomes (480, 640).**Q44**

A data analyst is computing statistics on a numeric dataset. Consider the following code using NumPy:

```
import numpy as np
data = np.array([1, 2, 3, 4, 5])
mean_val = np.mean(data)
var_val = np.var(data)
print(mean_val, var_val)
```

Question:

What output does this code produce?

A	3.0 2.0
B	3.0 2.5
C	3.5 2.5
D	3.5 3.0

Answer: A**Explanation:** The mean of [1, 2, 3, 4, 5] is 3.0, and the variance is calculated as the average of squared differences from the mean, which is 2.0.**Q45**

A student is developing a CNN for classifying images from a dataset like CIFAR-10. They construct a simple model using TensorFlow/Keras as shown:

```
import tensorflow as tf
from tensorflow.keras.models import Sequential
from tensorflow.keras.layers import Conv2D, MaxPooling2D, Flatten, Dense

model = Sequential([
    Conv2D(32, (3,3), activation='relu', input_shape=(32,32,3)),
    MaxPooling2D(pool_size=(2,2)),
```

```
Flatten(),
Dense(10, activation='softmax')
```

])

```
model.compile(optimizer='adam', loss='categorical_crossentropy')
model.summary()
```

Question:

What is the primary function of the MaxPooling2D layer in this CNN architecture?

CNN architecture: feature extraction to prediction



Pooling reduces spatial size and computation.

A	To increase the number of parameters for better learning.
B	To reduce spatial dimensions and control overfitting.
C	To convert the convolutional output into a one-dimensional vector.
D	To apply non-linear activation to the network.

Answer: B

Explanation: The MaxPooling2D layer reduces the spatial size of the representation, thereby decreasing the number of parameters and computation, and helping to control overfitting. It does not increase parameters, flatten the output, or apply activation functions.

Q46

A researcher is testing a convolutional layer on synthetic data. They use the following code:

```
import tensorflow as tf
import numpy as np

input_tensor = np.random.rand(1, 32, 32, 3)
conv_layer = tf.keras.layers.Conv2D(16, (3,3), activation='relu', padding='same')
output_tensor = conv_layer(input_tensor)
print(output_tensor.shape)
```

Question:

What will be the printed shape of output_tensor?

A	(1, 30, 30, 16)
B	(1, 32, 32, 16)
C	(1, 34, 34, 16)
D	(1, 32, 32, 3)

Answer: B

Explanation: With padding='same', the convolutional layer preserves the spatial dimensions. Thus, the output tensor retains a shape of (1, 32, 32, 16) (16 filters produce 16 channels).

Q47

A developer wants to create a decorator that logs the result of a function in uppercase. They implement the following code:

```
def uppercase_decorator(func):
    def wrapper(*args, **kwargs):
        original_result = func(*args, **kwargs)
        return original_result.upper()
    return wrapper

@uppercase_decorator
def greet(name):
    return f"hello {name}"

print(greet("Alice"))
```

Question:

What will this code print?

A	"hello Alice"
B	"HELLO ALICE"
C	"Hello Alice"
D	An error due to missing arguments in the decorator.

Answer: B

Explanation: The decorator uppercase_decorator modifies the output of greet by converting it to uppercase. Thus, the output is "HELLO ALICE".

Q48

A machine learning engineer is implementing a linear regression model to predict housing prices based on square footage. They use the following code:

```
import numpy as np
from sklearn.linear_model import LinearRegression

X = np.array([[800], [1000], [1200], [1500]])
y = np.array([150, 200, 250, 300])
model = LinearRegression()
model.fit(X, y)
print("Coefficient:", model.coef_[0], "Intercept:", model.intercept_)
```

Question:

What is the expected relationship learned by the model?

A	A negative coefficient with a large intercept
B	A positive coefficient indicating price increases with square footage
C	Zero coefficient implying no relationship
D	A complex non-linear relationship

Answer: B

Explanation: Given that housing prices typically increase with square footage, the linear regression model learns a positive coefficient, showing a direct relationship between square footage and price.

Q49

For a real-time AI project, a developer writes code to capture a single frame from a webcam. Consider the following code snippet:

```
import cv2
cap = cv2.VideoCapture(0)

ret, frame = cap.read() cap.release()

print(ret, frame.shape)
```

Question:

Assuming the webcam returns a color image with resolution 480×640, what is the expected output format?

A	True (480, 640)
----------	-----------------

B	True (480, 640, 3)
C	False None
D	True (640, 480, 3)

Answer: B

Explanation: If the frame is captured successfully, ret will be True and frame.shape will be (480, 640, 3), representing height, width, and the three color channels (BGR).

Q50

To improve the robustness of their CNN model, a data scientist employs data augmentation using Keras. They set up an image generator with rotation and zoom range:

```
from tensorflow.keras.preprocessing.image import ImageDataGenerator

datagen = ImageDataGenerator(
    rotation_range=40,
    zoom_range=0.2,
    horizontal_flip=True,
    fill_mode='nearest'
```

)

Question:

What is the primary purpose of specifying rotation_range=40 in the image generator?

A	To limit the image rotation to exactly 40 degrees.
B	To randomly rotate images by up to 40 degrees, enhancing model robustness to orientation variations.
C	To ensure all images are rotated by 40 degrees.
D	To convert images to grayscale after rotation.

Answer: B

Explanation: Setting rotation_range=40 allows the generator to randomly rotate images by up to 40 degrees in either direction during augmentation. This exposes the CNN to orientation variation and can improve generalization. It does not force every image to rotate exactly 40 degrees.

Answer Key Summary

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

