ش2	202110164 شهد مطلق راتب بركات_DS&AI
Started on	Thursday, 19 June 2025, 1:48 PM
State	Finished
Completed on	Thursday, 19 June 2025, 2:23 PM
Time taken	34 mins 46 secs
Grade	**17.00** out of 40.00 (**42.5**%)

Question 1

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Clarify the purpose of anchors in the Faster R-CNN architecture.

a. Reduce the model size

b. Store feature vectors ✕

c. Match pixel intensities

d. Represent potential object bounding boxes

The correct answer is: Represent potential object bounding boxes

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 13:51:15	Saved: Store feature vectors	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 2

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Choose the element that Mask R-CNN outputs for each detected instance.

a. Pixel-level masks
b. Anchor probabilities ✕
c. Color values
d. Aspect ratios

The correct answer is: Pixel-level masks

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 13:51:15	Saved: Anchor probabilities	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 3

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Combine two processes to describe the goal of object detection.

a. Classification and normalization
b. Segmentation and translation ✕
c. Alignment and resizing
d. Classification and localization

The correct answer is: Classification and localization

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 13:51:15	Saved: Segmentation and translation	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 4

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Describe how eye detection can be nested within face detection using Haar cascades.

- a. Apply feature merging to both classifiers ✕
- b. Train a single unified classifier
- c. Generate feature pyramids
- d. Use hierarchical detection where eyes are searched within face regions

The correct answer is: Use hierarchical detection where eyes are searched within face regions

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 13:52:54	Saved: Apply feature merging to both classifiers	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 5

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Recognize the advantage of using pre-trained Haar cascades in OpenCV.

a. Provides pixel-level segmentation masks
b. Requires large GPU memory ✕
c. Easily deployable for standard detection tasks
d. Detects identity of objects

The correct answer is: Easily deployable for standard detection tasks

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 13:52:54	Saved: Requires large GPU memory	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 6

Correct

Mark 1.00 out of 1.00

Hide sidebars

Choose the type of feature that convolutional layers typically extract from image data.

- a. Spatial features like edges and textures ✓
- b. Metadata
- c. Pixel coordinates
- d. Color labels

The correct answer is: Spatial features like edges and textures

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 13:52:54	Saved: Spatial features like edges and textures	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Correct**	**1.00**

Question 7

Correct

Mark 1.00 out of 1.00

Hide sidebars

Examine why R-CNN is computationally slower than Fast R-CNN.

a. Because region proposals are generated during inference
b. Because it includes semantic segmentation for each RoI
c. Because it uses deeper CNNs
d. Because it performs a forward pass for each region proposal individually ✓

The correct answer is: Because it performs a forward pass for each region proposal individually

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 13:54:49	Saved: Because it performs a forward pass for each region proposal individually	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Correct**	**1.00**

Question 8

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Determine how the RPN in Faster R-CNN predicts object locations.

- a. Using Gaussian filters ✕
- b. From thresholding segmentation masks
- c. By regressing bounding box coordinates for anchors
- d. With Softmax classification

The correct answer is: By regressing bounding box coordinates for anchors

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 13:54:49	Saved: Using Gaussian filters	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 9

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Identify the operation that Gaussian smoothing performs before edge detection.

- a. Reduces noise in the image
- b. Enhances edges
- c. Extracts corners ✕
- d. Computes gradients directly

The correct answer is: Reduces noise in the image

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 13:54:49	Saved: Extracts corners	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 10

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Determine the primary function of a pooling layer in a convolutional neural network.

a. Enhances edges
b. Converts grayscale to RGB
c. Reduces spatial dimensions
d. Applies softmax to outputs ✕

The correct answer is: Reduces spatial dimensions

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 13:57:23	Saved: Applies softmax to outputs	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 11

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Explain why Mask R-CNN separates class and mask predictions.

a. To improve bounding box regression
b. To allow parallel prediction of masks and classes
c. To reduce the number of classes ✕
d. To reuse weights

The correct answer is: To allow parallel prediction of masks and classes

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 13:57:23	Saved: To reduce the number of classes	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 12

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Point out the method R-CNN uses to generate region proposals.

- a. YOLO anchors ✕
- b. Region Proposal Network
- c. Attention maps
- d. Selective Search

The correct answer is: Selective Search

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 13:57:23	Saved: YOLO anchors	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 13

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Clarify how YOLOv8 distinguishes and counts objects by their class during real-time detection.

- a. By summing all bounding boxes ✕
- b. By filtering repeated frames
- c. By drawing all class labels
- d. By checking class IDs of each detection

The correct answer is: By checking class IDs of each detection

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:00:40	Saved: By summing all bounding boxes	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 14

Correct

Mark 1.00 out of 1.00

Hide sidebars

Describe what happens to each frame of a video when YOLOv8 is applied to it.

a. Grayscale conversion
b. Image resizing only
c. Gaussian blur
d. YOLOv8 inference and bounding box plotting ✓

The correct answer is: YOLOv8 inference and bounding box plotting

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:00:40	Saved: YOLOv8 inference and bounding box plotting	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Correct**	**1.00**

Question 15

Correct

Mark 1.00 out of 1.00

Hide sidebars

Within the Faster R-CNN pipeline, deduce why anchors of multiple aspect ratios are crucial.

a. They increase robustness to varying object shapes and scales ✓
b. They allow pixel-wise supervision of all background areas
c. They provide rotation invariance to the network
d. They reduce training time significantly

The correct answer is: They increase robustness to varying object shapes and scales

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:22:36	Saved: They increase robustness to varying object shapes and scales	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Correct**	**1.00**

Question 16

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Highlight the parameter in Haar detection that adjusts bounding box filtering sensitivity.

a. thresholdBinary ✕
b. minNeighbors
c. objectCount
d. maxROI

The correct answer is: minNeighbors

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:05:57	Saved: thresholdBinary	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 17

Correct

Mark 1.00 out of 1.00

Hide sidebars

Assess the architectural benefit of replacing the SVM classifier in R-CNN with a Softmax classifier in Fast R-CNN.

a. Removes the need for bounding box annotations
b. Permits end-to-end training of classification and regression ✓
c. Enables detection of object orientation
d. Supports non-Euclidean feature spaces

The correct answer is: Permits end-to-end training of classification and regression

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:05:57	Saved: Permits end-to-end training of classification and regression	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Correct**	**1.00**

Question 18

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Choose the model type suitable for performing instance segmentation in YOLOv8.

a. yolov8s-seg.pt

b. yolov8s-cls-segment.pt

c. yolov8n.pt

d. yolov8m-det.pt ✕

The correct answer is: yolov8s-seg.pt

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:05:57	Saved: yolov8m-det.pt	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 19

Correct

Mark 1.00 out of 1.00

Hide sidebars

Explain how Mask R-CNN achieves class-specific segmentation without competition.

a. Each class is represented by its own color code in the output
b. The mask prediction is post-processed using CRF to isolate classes
c. Each RoI has separate mask outputs for each class, avoiding softmax ✓
d. A single mask branch generates multiple masks with shared weights

The correct answer is: Each RoI has separate mask outputs for each class, avoiding softmax

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:10:11	Saved: Each RoI has separate mask outputs for each class, avoiding softmax	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Correct**	**1.00**

Question 20

Correct

Mark 1.00 out of 1.00

Hide sidebars

Identify the optimizer often used for CNNs due to its adaptive learning rate.

a. Momentum
b. Adam ✓
c. RMSprop
d. Adagrad

The correct answer is: Adam

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:10:11	Saved: Adam	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Correct**	**1.00**

Question 21

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Reveal the main reason for applying geometric verification after matching feature descriptors.

- a. To visualize matches with bounding boxes ✕
- b. To convert grayscale images to RGB
- c. To remove false matches using spatial constraints
- d. To count detected objects per frame

The correct answer is: To remove false matches using spatial constraints

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:10:11	Saved: To visualize matches with bounding boxes	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 22

Correct

Mark 1.00 out of 1.00

Hide sidebars

Determine why hysteresis is used in the Canny edge detection algorithm.

a. To suppress all low gradients
b. To quantize orientations to 90 degrees
c. To avoid high-frequency filtering
d. To connect weak edges based on strong edge support ✓

The correct answer is: To connect weak edges based on strong edge support

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:13:00	Saved: To connect weak edges based on strong edge support	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Correct**	**1.00**

Question 23

Correct

Mark 1.00 out of 1.00

Hide sidebars

Evaluate the main reason Mask R-CNN adds a third branch to the Faster R-CNN structure.

a. To increase the detection threshold

b. To perform instance segmentation with per-pixel predictions ✓

c. To allow probabilistic classification with dropout

d. To optimize with adversarial training

The correct answer is: To perform instance segmentation with per-pixel predictions

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:13:00	Saved: To perform instance segmentation with per-pixel predictions	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Correct**	**1.00**

Question 24

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

During training of Mask R-CNN, identify the role of the K-class mask prediction per RoI.

a. It ensures every mask shares the same weights across classes
b. It merges overlapping object proposals using non-maximum suppression
c. It forces binary segmentation for all background objects ✕
d. It allows mask generation to be decoupled from class competition

The correct answer is: It allows mask generation to be decoupled from class competition

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:13:00	Saved: It forces binary segmentation for all background objects	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 25

Correct

Mark 1.00 out of 1.00

Hide sidebars

Identify the best way to iterate over all image files in a folder when applying YOLOv8.

- a. listdir loop
- b. os.path.join
- c. glob.glob inside a for loop ✓
- d. pandas loop

The correct answer is: glob.glob inside a for loop

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:14:56	Saved: glob.glob inside a for loop	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Correct**	**1.00**

Question 26

Correct

Mark 1.00 out of 1.00

Hide sidebars

In the context of object detection, determine why bounding box regression uses Smooth L1 loss instead of L2.

a. Smooth L1 directly minimizes Intersection over Union

b. L2 loss increases numerical stability

c. Smooth L1 is more robust to outliers and avoids exploding gradients ✓

d. L2 better adapts to imbalanced datasets

The correct answer is: Smooth L1 is more robust to outliers and avoids exploding gradients

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:14:56	Saved: Smooth L1 is more robust to outliers and avoids exploding gradients	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Correct**	**1.00**

Question 27

Correct

Mark 1.00 out of 1.00

Hide sidebars

Recognize the module introduced by Faster R-CNN to eliminate Selective Search.

a. Segmentation Head
b. Anchor Pooling
c. Region Proposal Network (RPN) ✓
d. Region Classifier

The correct answer is: Region Proposal Network (RPN)

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:14:56	Saved: Region Proposal Network (RPN)	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Correct**	**1.00**

Question 28

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Select the part of a Haar cascade used for detecting candidate regions.

a. Deep feature encoders
b. Keypoint regression head
c. Sliding window classifier
d. Gaussian filters ✕

The correct answer is: Sliding window classifier

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:17:00	Saved: Gaussian filters	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 29

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Canny edge detection involves non-maximum suppression for what main purpose?

a. Thin edge responses to precise locations

b. Increase computation speed ✕

c. Connect broken edges

d. Segment regions with high intensity

The correct answer is: Thin edge responses to precise locations

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:17:00	Saved: Increase computation speed	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 30

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Choose the behavior exhibited by a MaxPooling2D operation.

a. Summing all pixels in a region
b. Selecting the maximum value from each kernel region
c. Removing padding from convolution ✕
d. Selecting the smallest value

The correct answer is: Selecting the maximum value from each kernel region

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:17:00	Saved: Removing padding from convolution	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 31

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Explain the effect of using bilinear interpolation in RoIAlign instead of quantization.

a. Applies a fixed grid to resize features uniformly
b. Improves classification but worsens mask quality
c. Increases overall training time without affecting accuracy ✕
d. Allows sub-pixel accuracy in feature extraction, improving segmentation

The correct answer is: Allows sub-pixel accuracy in feature extraction, improving segmentation

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:18:59	Saved: Increases overall training time without affecting accuracy	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 32

Correct

Mark 1.00 out of 1.00

Hide sidebars

Determine the visual result of executing `results[0].plot()` in a YOLOv8 detection workflow.

a. Image tensor matrix
b. Histogram of class probabilities
c. Image with bounding boxes and labels ✓
d. Confidence heatmap

The correct answer is: Image with bounding boxes and labels

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:18:59	Saved: Image with bounding boxes and labels	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Correct**	**1.00**

Question 33

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

In Faster R-CNN, determine how the RPN handles overlapping anchor predictions.

a. Using a dense grid where each pixel becomes an anchor
b. Through non-maximum suppression to retain high-scoring distinct proposals
c. By averaging coordinates of all overlapping anchors
d. By converting anchors to bounding boxes using confidence score thresholds ✕

The correct answer is: Through non-maximum suppression to retain high-scoring distinct proposals

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:18:59	Saved: By converting anchors to bounding boxes using confidence score thresholds	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 34

Correct

Mark 1.00 out of 1.00

Hide sidebars

Identify the type of output generated by the mask branch in Mask R-CNN.

a. Class labels only
b. Bounding box center coordinates
c. Color histograms
d. Binary segmentation maps ✓

The correct answer is: Binary segmentation maps

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:20:42	Saved: Binary segmentation maps	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Correct**	**1.00**

Question 35

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

After defining a CNN model architecture, select the next step before training.

a. Plotting accuracy
b. Compiling the model with loss and optimizer
c. Collecting images ✕
d. Testing the model

The correct answer is: Compiling the model with loss and optimizer

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:20:42	Saved: Collecting images	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 36

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

What the improvement Faster R-CNN brings compared to Fast R-CNN.

a. Adds ReLU activation ✕

b. Removes background class

c. Converts boxes to masks

d. Learns proposals via RPN instead of Selective Search

The correct answer is: Learns proposals via RPN instead of Selective Search

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:20:42	Saved: Adds ReLU activation	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 37

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Reveal the reason why Fast R-CNN performs faster than its predecessor.

a. Single CNN pass with RoI pooling

b. Binary image pre-processing ✕

c. GPU requirement

d. Use of deeper networks

The correct answer is: Single CNN pass with RoI pooling

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:21:52	Saved: Binary image pre-processing	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Incorrect**	**0.00**

Question 38

Correct

Mark 1.00 out of 1.00

Hide sidebars

Select the justification for applying feature map sharing in Fast R-CNN.

a. It reduces feature dimensionality by using average pooling

b. It performs anchor refinement before classification

c. It avoids redundant computation by reusing convolution outputs for all regions ✓

d. It prevents overfitting during testing

The correct answer is: It avoids redundant computation by reusing convolution outputs for all regions

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:21:52	Saved: It avoids redundant computation by reusing convolution outputs for all regions	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Correct**	**1.00**

Question 39

Correct

Mark 1.00 out of 1.00

Hide sidebars

Explain how the classification mode in YOLOv8 differs from object detection mode.

a. It uses region proposal networks
b. It outputs bounding boxes only
c. It assigns a label to the entire image without bounding boxes ✓
d. It detects masks for each object

The correct answer is: It assigns a label to the entire image without bounding boxes

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:21:52	Saved: It assigns a label to the entire image without bounding boxes	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Correct**	**1.00**

Question 40

Correct

Mark 1.00 out of 1.00

Hide sidebars

Explain the benefit of RoI Pooling in Fast R-CNN.

a. Resize input images
b. Generate anchors
c. Normalize feature maps
d. Extract fixed-size feature vectors from proposals ✓

The correct answer is: Extract fixed-size feature vectors from proposals

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:26	Started	Not yet answered	
2	19/06/25, 14:22:05	Saved: Extract fixed-size feature vectors from proposals	Answer saved	
3	**19/06/25, 14:23:12**	**Attempt finished**	**Correct**	**1.00**

◀ Assignment 5 marks

Jump to...

Table of contents

- General
- Topic 1
- Topic 2
- Topic 3
- Topic 4
- Topic 5
- Topic 6
- Topic 7
- Topic 8
- Topic 9
- Topic 10