



[E-Learning Portal](#)
[Resources](#)
[إدماج التعلم الإلكتروني](#)
[أيقونة التعلم الإلكتروني](#)
[Old Moodle](#)
[FIT-Courses](#)
[Quality](#)

My courses (5) ▾



الرؤية الحاسوبية (1_606384)

[Dashboard](#)
[My courses](#)
[606384_1](#)
[Topic 18](#)
[Final Exam](#)

[Dashboard](#)
[Site home](#)
[Calendar](#)
[Badges](#)
[All courses](#)
[Edit course settings](#)
[Content bank](#)
[Turn editing on](#)

Course management

ت2 [DS&AI](#) [تميم محمد انيس صالح صالح 202210369](#)

Started on Thursday, 19 June 2025, 1:54 PM

State Finished

Completed on Thursday, 19 June 2025, 2:20 PM

Time taken 25 mins 46 secs

Grade 19.00 out of 40.00 (47.5%)

Question 1

Incorrect

Mark 0.00 out of
1.00

Hide sidebars

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

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

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

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 13:56:05	Saved: They reduce training time significantly	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

Question 2

Incorrect

Mark 0.00 out of
1.00

Hide sidebars

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

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

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

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 13:56:05	Saved: To visualize matches with bounding boxes	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

Question 3

Incorrect

Mark 0.00 out of
1.00

Hide sidebars

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

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

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, 13:54:21	Started	Not yet answered	
2	19/06/25, 13:56:05	Saved: Each RoI has separate mask outputs for each class, avoiding softmax	Answer saved	
3	19/06/25, 14:07:40	Saved: A single mask branch generates multiple masks with shared weights	Answer saved	
4	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

Question 4

Correct

Mark 1.00 out of 1.00

Hide sidebars

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

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

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

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 13:57:36	Saved: Extract fixed-size feature vectors from proposals	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Correct	1.00

Question 5

Incorrect

Mark 0.00 out of
1.00

Hide sidebars

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

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

The correct answer is: Easily deployable for standard detection tasks

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 13:57:36	Saved: Detects identity of objects	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

Question 6

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

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, 13:54:21	Started	Not yet answered	
2	19/06/25, 13:57:36	Saved: It forces binary segmentation for all background objects	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

Question 7

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

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

The correct answer is: Pixel-level masks

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:00:01	Saved: Pixel-level masks	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Correct	1.00

Question 8

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. Image resizing only
- b. Gaussian blur
- c. Grayscale conversion
- 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, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:00:01	Saved: YOLOv8 inference and bounding box plotting	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Correct	1.00

Question 9

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

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

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

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:00:01	Saved: Allows sub-pixel accuracy in feature extraction, improving segmentation	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Correct	1.00

Question 10

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

The correct answer is: Image with bounding boxes and labels

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:02:30	Saved: Image with bounding boxes and labels	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Correct	1.00

Question 11

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. By averaging coordinates of all overlapping anchors
- c. Through non-maximum suppression to retain high-scoring distinct proposals
- 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, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:02:30	Saved: By converting anchors to bounding boxes using confidence score thresholds	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

Question 12

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

- a. Learns proposals via RPN instead of Selective Search ✓
- b. Adds ReLU activation
- c. Converts boxes to masks
- d. Removes background class

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

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:02:30	Saved: Learns proposals via RPN instead of Selective Search	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Correct	1.00

Question 13

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 connect weak edges based on strong edge support ✓
- c. To avoid high-frequency filtering
- d. To quantize orientations to 90 degrees

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

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:04:29	Saved: To connect weak edges based on strong edge support	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Correct	1.00

Question 14

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 uses deeper CNNs
- c. Because it includes semantic segmentation for each RoI
- 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, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:04:29	Saved: Because it performs a forward pass for each region proposal individually	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Correct	1.00

Question 15

Correct

Mark 1.00 out of 1.00

Hide sidebars

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

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

The correct answer is: Selective Search

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:04:29	Saved: Selective Search	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Correct	1.00

Question 16

Incorrect

Mark 0.00 out of
1.00

Hide sidebars

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

- a. Binary image pre-processing ✘
- b. GPU requirement
- c. Use of deeper networks
- d. Single CNN pass with RoI pooling

The correct answer is: Single CNN pass with RoI pooling

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:06:40	Saved: Binary image pre-processing	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

Question 17

Incorrect

Mark 0.00 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. L2 better adapts to imbalanced datasets
- b. Smooth L1 directly minimizes Intersection over Union ✖
- c. Smooth L1 is more robust to outliers and avoids exploding gradients
- d. L2 loss increases numerical stability

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, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:06:40	Saved: Smooth L1 directly minimizes Intersection over Union	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

Question 18

Incorrect

Mark 0.00 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. Enables detection of object orientation ✖
- b. Supports non-Euclidean feature spaces
- c. Removes the need for bounding box annotations
- d. Permits end-to-end training of classification and regression

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

Response history

Step	Time	Action	State	Marks
<u>1</u>	19/06/25, 13:54:21	Started	Not yet answered	
<u>2</u>	19/06/25, 14:06:40	Saved: Enables detection of object orientation	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

Question 19

Incorrect

Mark 0.00 out of
1.00

Hide sidebars

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

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

The correct answer is: Reduces spatial dimensions

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:09:42	Saved: Applies softmax to outputs	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

Question **20**

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

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

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, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:09:42	Saved: Use hierarchical detection where eyes are searched within face regions	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Correct	1.00

Question 21

Incorrect

Mark 0.00 out of
1.00

Hide sidebars

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

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

The correct answer is: Sliding window classifier

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:09:42	Saved: Gaussian filters	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

Question 22

Incorrect

Mark 0.00 out of
1.00

Hide sidebars

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

- a. Momentum ✘
- b. Adagrad
- c. Adam
- d. RMSprop

The correct answer is: Adam

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:12:40	Saved: Momentum	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

Question **23**

Incorrect

Mark 0.00 out of
1.00

Hide sidebars

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

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

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

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:12:40	Saved: Collecting images	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

Question 24

Incorrect

Mark 0.00 out of
1.00

Hide sidebars

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

- a. objectCount
- b. thresholdBinary ✖
- c. maxROI
- d. minNeighbors

The correct answer is: minNeighbors

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:12:40	Saved: thresholdBinary	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

Question 25

Incorrect

Mark 0.00 out of
1.00

Hide sidebars

Identify the operation that Gaussian smoothing performs before edge detection.

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

The correct answer is: Reduces noise in the image

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:19:59	Saved: Enhances edges	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

Question **26**

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

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

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

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:19:59	Saved: By regressing bounding box coordinates for anchors	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Correct	1.00

Question 27

Incorrect

Mark 0.00 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 prevents overfitting during testing ✘
- d. It avoids redundant computation by reusing convolution outputs for all regions

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, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:19:08	Saved: It prevents overfitting during testing	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

Question 28

Correct

Mark 1.00 out of 1.00

Hide sidebars

Choose the behavior exhibited by a MaxPooling2D operation.

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

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

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:18:32	Saved: Selecting the maximum value from each kernel region	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Correct	1.00

Question 29

Correct

Mark 1.00 out of 1.00

Hide sidebars

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

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

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

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:18:32	Saved: To allow parallel prediction of masks and classes	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Correct	1.00

Question 30

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 optimize with adversarial training
- c. To allow probabilistic classification with dropout
- d. To perform instance segmentation with per-pixel predictions ✓

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

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:18:32	Saved: To perform instance segmentation with per-pixel predictions	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Correct	1.00

Question 31

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

- a. yolov8n.pt
- b. yolov8s-cls-segment.pt
- c. yolov8m-det.pt
- d. yolov8s-seg.pt ✓

The correct answer is: yolov8s-seg.pt

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:17:53	Saved: yolov8s-seg.pt	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Correct	1.00

Question 32

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

The correct answer is: Binary segmentation maps

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:17:53	Saved: Binary segmentation maps	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Correct	1.00

Question 33

Correct

Mark 1.00 out of 1.00

Hide sidebars

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

- a. Reduce the model size
- b. Match pixel intensities
- c. Store feature vectors
- 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, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:17:53	Saved: Represent potential object bounding boxes	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Correct	1.00

Question **34**

Incorrect

Mark 0.00 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. pandas loop ✘
- c. glob.glob inside a for loop
- d. os.path.join

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

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:17:11	Saved: pandas loop	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

Question 35

Incorrect

Mark 0.00 out of
1.00

Hide sidebars

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

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

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, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:17:11	Saved: It detects masks for each object	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

Question 36

Incorrect

Mark 0.00 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. Color labels
- c. Metadata
- d. Pixel coordinates ✖

The correct answer is: Spatial features like edges and textures

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:17:11	Saved: Pixel coordinates	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

Question 37

Incorrect

Mark 0.00 out of
1.00

Hide sidebars

Combine two processes to describe the goal of object detection.

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

The correct answer is: Classification and localization

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:14:12	Saved: Classification and normalization	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

Question 38

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. Region Classifier
- c. Anchor Pooling
- d. Region Proposal Network (RPN) ✓

The correct answer is: Region Proposal Network (RPN)

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:14:12	Saved: Region Proposal Network (RPN)	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Correct	1.00

Question 39

Correct

Mark 1.00 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 drawing all class labels
- c. By filtering repeated frames
- 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, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:14:12	Saved: By checking class IDs of each detection	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Correct	1.00

Question 40

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

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

- a. Increase computation speed
- b. Thin edge responses to precise locations
- 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, 13:54:21	Started	Not yet answered	
2	19/06/25, 14:13:01	Saved: Segment regions with high intensity	Answer saved	
3	19/06/25, 14:20:07	Attempt finished	Incorrect	0.00

◀ 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

Hide sidebars

> **Topic 11**

> **Topic 12**

> **Topic 13**

> **Topic 14**

> **Topic 15**

> **Topic 16**

> **Topic 17**

✓ **Topic 18**

[yolo](#)

[yolo\(1\)](#)

[CNN TUTORIAL](#)

[CNN CODE 2](#)

[r cnn CODE](#)

[quiz](#) 

[Project Submission 10-6-2025](#)

[Assignment 5 marks](#)

[Final Exam](#) 

Copyright © PUeLC. 2024-2025 All rights reserved.



You are logged in as [د. حسام نواف سليمان فاخوري](#) (Log out) [Help and documentation](#)

[Data retention summary](#)