



[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

21 [202111335](#) [إبراهيم احمد جابر الطيطي](#) [DS&AI](#)

Started on Thursday, 19 June 2025, 1:48 PM

State Finished

Completed on Thursday, 19 June 2025, 2:44 PM

Time taken 56 mins 16 secs

Grade 37.00 out of 40.00 (92.5%)

Question 1

Correct

Mark 1.00 out of 1.00

Hide sidebars

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

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

The correct answer is: Easily deployable for standard detection tasks

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:50:21	Saved: Easily deployable for standard detection tasks	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 2

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

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

The correct answer is: Spatial features like edges and textures

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:50:21	Saved: Spatial features like edges and textures	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 3

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

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

The correct answer is: Reduces spatial dimensions

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:50:21	Saved: Reduces spatial dimensions	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 4

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

- a. Generate feature pyramids
- b. Train a single unified classifier
- c. Apply feature merging to both classifiers
- 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, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:53:44	Saved: Use hierarchical detection where eyes are searched within face regions	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 5

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

- a. To avoid high-frequency filtering
- b. To connect weak edges based on strong edge support ✓
- c. To suppress all low gradients
- 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, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:53:44	Saved: To connect weak edges based on strong edge support	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 6

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

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

The correct answer is: Selective Search

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:53:44	Saved: Selective Search	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 7

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

- a. Represent potential object bounding boxes ✓
- b. Store feature vectors
- c. Reduce the model size
- d. Match pixel intensities

The correct answer is: Represent potential object bounding boxes

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:55:02	Saved: Represent potential object bounding boxes	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 8

Correct

Mark 1.00 out of
1.00

Hide sidebars

Combine two processes to describe the goal of object detection.

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

The correct answer is: Classification and localization

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:55:02	Saved: Classification and localization	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 9

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

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

The correct answer is: Region Proposal Network (RPN)

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:55:02	Saved: Region Proposal Network (RPN)	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 10

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

- a. It assigns a label to the entire image without bounding boxes ✓
- b. It outputs bounding boxes only
- c. It uses region proposal networks
- 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, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:54:42	Saved: It assigns a label to the entire image without bounding boxes	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 11

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

- a. Compiling the model with loss and optimizer ✓
- b. Collecting images
- c. Plotting accuracy
- 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, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:54:42	Saved: Compiling the model with loss and optimizer	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question **12**

Correct

Mark 1.00 out of 1.00

Hide sidebars

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

- a. To allow parallel prediction of masks and classes ✓
- b. To improve bounding box regression
- 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, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:54:42	Saved: To allow parallel prediction of masks and classes	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 13

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

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

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:55:21	Saved: To perform instance segmentation with per-pixel predictions	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 14

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. Normalize feature maps
- c. Resize input images
- d. Generate anchors

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

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:55:21	Saved: Extract fixed-size feature vectors from proposals	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 15

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. Increases overall training time without affecting accuracy
- d. Applies a fixed grid to resize features uniformly

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

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:55:21	Saved: Allows sub-pixel accuracy in feature extraction, improving segmentation	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question **16**

Correct

Mark 1.00 out of 1.00

Hide sidebars

Identify the operation that Gaussian smoothing performs before edge detection.

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

The correct answer is: Reduces noise in the image

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:58:55	Saved: Reduces noise in the image	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 17

Incorrect

Mark 0.00 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. Class labels only
- d. Binary segmentation maps

The correct answer is: Binary segmentation maps

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:58:55	Saved: Color histograms	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Incorrect	0.00

Question 18

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, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:58:55	Saved: YOLOv8 inference and bounding box plotting	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 19

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

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

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

The correct answer is: yolov8s-seg.pt

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:58:20	Saved: yolov8s-cls-segment.pt	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Incorrect	0.00

Question 20

Correct

Mark 1.00 out of 1.00

Hide sidebars

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

- a. To remove false matches using spatial constraints ✓
- b. To convert grayscale images to RGB
- c. To visualize matches with bounding boxes
- 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, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:58:20	Saved: To remove false matches using spatial constraints	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question **21**

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

- a. It performs anchor refinement before classification
- b. It avoids redundant computation by reusing convolution outputs for all regions ✓
- c. It reduces feature dimensionality by using average pooling
- 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, 13:48:05	Started	Not yet answered	
2	19/06/25, 13:58:20	Saved: It avoids redundant computation by reusing convolution outputs for all regions	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question **22**

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

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, 13:48:05	Started	Not yet answered	
<u>2</u>	19/06/25, 14:02:07	Saved: Permits end-to-end training of classification and regression	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question **23**

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

- a. Adds ReLU activation
- b. Converts boxes to masks
- c. Removes background class
- 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, 13:48:05	Started	Not yet answered	
2	19/06/25, 14:02:07	Saved: Learns proposals via RPN instead of Selective Search	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question **24**

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. Each RoI has separate mask outputs for each class, avoiding softmax ✓
- c. A single mask branch generates multiple masks with shared weights
- d. The mask prediction is post-processed using CRF to isolate classes

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, 13:48:05	Started	Not yet answered	
2	19/06/25, 14:02:07	Saved: Each RoI has separate mask outputs for each class, avoiding softmax	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 25

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

The correct answer is: Image with bounding boxes and labels

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 14:03:01	Saved: Image with bounding boxes and labels	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question **26**

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

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

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

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 14:03:01	Saved: By regressing bounding box coordinates for anchors	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 27

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 reduce training time significantly
- c. They allow pixel-wise supervision of all background areas
- d. They provide rotation invariance to the network

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

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 14:03:01	Saved: They increase robustness to varying object shapes and scales	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question **28**

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

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

The correct answer is: Single CNN pass with RoI pooling

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 14:05:09	Saved: Single CNN pass with RoI pooling	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 29

Correct

Mark 1.00 out of 1.00

Hide sidebars

Choose the behavior exhibited by a MaxPooling2D operation.

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

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

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 14:05:09	Saved: Selecting the maximum value from each kernel region	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 30

Correct

Mark 1.00 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. Gaussian filters
- d. Sliding window classifier ✓

The correct answer is: Sliding window classifier

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 14:05:09	Saved: Sliding window classifier	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 31

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

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

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, 13:48:05	Started	Not yet answered	
2	19/06/25, 14:05:23	Saved: Because it performs a forward pass for each region proposal individually	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 32

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

- a. Connect broken edges
- b. Segment regions with high intensity
- c. Thin edge responses to precise locations ✓
- d. Increase computation speed

The correct answer is: Thin edge responses to precise locations

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 14:05:23	Saved: Thin edge responses to precise locations	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 33

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 better adapts to imbalanced datasets
- 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, 13:48:05	Started	Not yet answered	
2	19/06/25, 14:05:23	Saved: Smooth L1 is more robust to outliers and avoids exploding gradients	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 34

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

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

The correct answer is: minNeighbors

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 14:06:24	Saved: minNeighbors	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 35

Correct

Mark 1.00 out of
1.00

Hide sidebars

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

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

The correct answer is: Pixel-level masks

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 14:06:24	Saved: Anchor probabilities	Answer saved	
3	19/06/25, 14:07:20	Saved: Pixel-level masks	Answer saved	
4	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 36

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

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

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 14:06:24	Saved: By summing all bounding boxes	Answer saved	
3	19/06/25, 14:06:52	Saved: By checking class IDs of each detection	Answer saved	
4	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 37

Incorrect

Mark 0.00 out of
1.00

Hide sidebars

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

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

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, 13:48:05	Started	Not yet answered	
2	19/06/25, 14:06:48	Saved: By converting anchors to bounding boxes using confidence score thresholds	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Incorrect	0.00

Question 38

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

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

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 14:07:35	Saved: <code>glob.glob</code> inside a for loop	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 39

Correct

Mark 1.00 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 merges overlapping object proposals using non-maximum suppression
- b. It ensures every mask shares the same weights across classes
- c. It allows mask generation to be decoupled from class competition ✓
- d. It forces binary segmentation for all background objects

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, 13:48:05	Started	Not yet answered	
2	19/06/25, 14:07:35	Saved: It allows mask generation to be decoupled from class competition	Answer saved	
3	19/06/25, 14:44:21	Attempt finished	Correct	1.00

Question 40

Correct

Mark 1.00 out of 1.00

Hide sidebars

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

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

The correct answer is: Adam

Response history

Step	Time	Action	State	Marks
1	19/06/25, 13:48:05	Started	Not yet answered	
2	19/06/25, 14:07:57	Saved: Adam	Answer saved	
3	19/06/25, 14:44:21	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

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](#)