



[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 15](#) > [Midterm Exam](#)

Dashboard | 
 Site home | 
 Calendar | 
 Badges | 
 All courses | 
 Edit course settings | 
 Content bank | 
 Turn editing on

Course management

ش2 [DS&AI](#) [شاهد مطلق راتب بركات 202110164](#)

**Started on** Tuesday, 6 May 2025, 12:02 PM

**State** Finished

**Completed on** Tuesday, 6 May 2025, 12:42 PM

**Time taken** 39 mins 48 secs

**Marks** 8.00/32.00

**Grade** 7.50 out of 30.00 (25%)

## Question 1

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q15. A batch size of 64 during training implies

- a. Accuracy is computed after every 64 epochs
- b. The model trains for only 64 iterations ✖
- c. Weights are updated every 64 gradient calculations
- d. Each epoch uses 64 samples in total

The correct answer is: Weights are updated every 64 gradient calculations

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:33:33	Saved: The model trains for only 64 iterations	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

## Question 2

Correct

Mark 1.00 out of  
1.00

Hide sidebars

Q14. Filters in CNNs are learned using

- a. Principal component analysis
- b. Histogram equalization
- c. Gradient descent and backpropagation over convolutional layers ✓
- d. Gaussian mixture models

The correct answer is: Gradient descent and backpropagation over convolutional layers

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:04:34	Saved: Gradient descent and backpropagation over convolutional layers	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Correct	1.00

## Question 3

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q1. Describing interest points using feature vectors helps in

- a. Compressing image resolution for lossy storage
- b. Matching semantically similar regions across images
- c. Reducing color intensity variation ✖
- d. Tracking object deformation over time

The correct answer is: Matching semantically similar regions across images

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:05:21	Saved: Reducing color intensity variation	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

Question 4

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Q6. Feature detection in the human visual system is functionally analogous to

- a. Lossless encoding in image formats
- b. Convolutional layers in deep networks
- c. Thresholding techniques in segmentation ✘
- d. JPEG compression in imaging devices

The correct answer is: Convolutional layers in deep networks

Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:35:17	Saved: Thresholding techniques in segmentation	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

## Question 5

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q14. Boundary detection failures typically arise when

- a. Optical zoom is reduced
- b. High frequency filters are applied ✘
- c. Local contrast is insufficient
- d. Color saturation is too low

The correct answer is: Local contrast is insufficient

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:07:05	Saved: High frequency filters are applied	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

## Question 6

Correct

Mark 1.00 out of  
1.00

Hide sidebars

Q19. A fully connected Dense layer with ReLU activation introduces

- a. Spatial invariance across filters
- b. Gradient clipping across pooling regions
- c. Nonlinear transformations after convolution ✓
- d. Reduced model depth for faster training

The correct answer is: Nonlinear transformations after convolution

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:07:48	Saved: Nonlinear transformations after convolution	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Correct	1.00

## Question 7

Correct

Mark 1.00 out of  
1.00

Hide sidebars

Q19. Image colorization tasks require

- a. Automatic dilation of binary shapes
- b. Fourier reconstruction of missing edge profiles
- c. Chain code approximation of image masks
- d. Filling grayscale regions with plausible chromatic values ✓

The correct answer is: Filling grayscale regions with plausible chromatic values

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:08:34	Saved: Filling grayscale regions with plausible chromatic values	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Correct	1.00

## Question 8

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q5. A 5x5x3 filter applied to a 32x32x3 input produces

- a. A single scalar per input image
- b. A 64x64 output using zero padding
- c. A 3D output with reduced spatial resolution ✖
- d. 75-dimensional dot product results per location

The correct answer is: 75-dimensional dot product results per location

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:09:48	Saved: A 64x64 output using zero padding	Answer saved	
3	6/05/25, 12:37:29	Saved: A 3D output with reduced spatial resolution	Answer saved	
4	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

## Question 9

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q1. CNNs excel in image analysis because they

- a. Operate on raw pixel intensities without transformation
- b. Convert images to binary maps directly ✖
- c. Use spatial hierarchies to extract multi-scale features
- d. Apply gradient filters for text-based input

The correct answer is: Use spatial hierarchies to extract multi-scale features

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:11:46	Saved: Use spatial hierarchies to extract multi-scale features	Answer saved	
3	6/05/25, 12:38:07	Saved: Convert images to binary maps directly	Answer saved	
4	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

Question 10

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Q8. Otsu's thresholding automatically

- a. Adjusts RGB channels based on luminance ✘
- b. Applies a fixed binary threshold
- c. Maximizes inter-class variance for threshold separation
- d. Minimizes pixel intensity values globally

The correct answer is: Maximizes inter-class variance for threshold separation

### Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:13:04	Saved: Adjusts RGB channels based on luminance	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

## Question 11

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q16. The speed advantage of Faster R-CNN arises because

- a. It uses binary SVMs instead of softmax
- b. It filters regions using PCA
- c. Convolution is done once and reused across tasks
- d. It removes fully connected layers from detection heads ✖

The correct answer is: Convolution is done once and reused across tasks

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:13:58	Saved: It removes fully connected layers from detection heads	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

## Question 12

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q5. The use of 28x28x1 shape in CNN input ensures that

- a. Feature maps use dropout as regularization ✘
- b. Channel depth is preserved for grayscale images
- c. Image patches are square and color coded
- d. Images are flattened for RNN input

The correct answer is: Channel depth is preserved for grayscale images

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:39:52	Saved: Feature maps use dropout as regularization	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

## Question 13

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q1. Edge detection in image processing aims to extract

- a. RGB values with reduced noise ✘
- b. Boundary-like features that encode shape
- c. Uniform pixel intensity values
- d. Contours generated from thresholding color

The correct answer is: Boundary-like features that encode shape

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:15:44	Saved: RGB values with reduced noise	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

## Question 14

Correct

Mark 1.00 out of  
1.00

Hide sidebars

Q11. Threshold ratio in edge hysteresis is typically designed to

- a. Balance edge sensitivity and noise tolerance ✓
- b. Ensure equal gradient division across axes
- c. Eliminate weak pixel structures altogether
- d. Match histogram ranges in smoothed images

The correct answer is: Balance edge sensitivity and noise tolerance

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:16:26	Saved: Balance edge sensitivity and noise tolerance	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Correct	1.00

## Question 15

Correct

Mark 1.00 out of  
1.00

Hide sidebars

Q17. The `Rectangle` function in matplotlib is used to

- a. Create a visual box around matched areas ✓
- b. Annotate noise in grayscale gradients
- c. Transform patches into binary
- d. Replace edge maps with histograms

The correct answer is: Create a visual box around matched areas

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:16:51	Saved: Create a visual box around matched areas	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Correct	1.00

Question **16**

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q6. Pooling layers in CNNs are used to

- a. Reduce dimensionality while retaining key features
- b. Replace convolution layers in shallow models
- c. Increase the number of filters at each level
- d. Normalize data for better color contrast **×**

The correct answer is: Reduce dimensionality while retaining key features

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:18:09	Saved: Normalize data for better color contrast	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

## Question 17

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q10. Haar Cascades cannot identify specific object instances because they

- a. Use unsupervised clustering of pixel data
- b. Only recognize shape templates, not identity
- c. Lack capability for multi-class classification ✖
- d. Rely on LBP rather than CNNs

The correct answer is: Only recognize shape templates, not identity

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:21:12	Saved: Lack capability for multi-class classification	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

## Question 18

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q5. Gaussian smoothing precedes derivative estimation to

- a. Introduce frequency-domain constraints
- b. Remove chromatic aberrations
- c. Limit dynamic range in gradient magnitude ✖
- d. Suppress high-frequency noise that misleads gradient computation

The correct answer is: Suppress high-frequency noise that misleads gradient computation

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:41:18	Saved: Limit dynamic range in gradient magnitude	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

## Question 19

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q5. Dilation in binary images leads to

- a. Enlargement of white regions, closing small black holes ✓
- b. Skewness correction of foreground objects
- c. Intensity normalization across object boundaries
- d. Subdivision of connected components

The correct answer is: Enlargement of white regions, closing small black holes

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:41:36	Saved: Enlargement of white regions, closing small black holes	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Correct	1.00

## Question 20

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q11. The `cv2.cornerHarris()` function requires image conversion to

- a. uint8
- b. float64
- c. uint16
- d. float32 ✓

The correct answer is: float32

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:23:46	Saved: float32	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Correct	1.00

## Question 21

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q3. Digital image representation stores pixel intensities

- a. In floating point arrays between -1 and 1
- b. Using Fourier basis functions
- c. In discrete grid positions with intensity values
- d. As RGB hex codes for compression ✖

The correct answer is: In discrete grid positions with intensity values

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:24:21	Saved: As RGB hex codes for compression	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

Question **22**

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q12. Biometrics often depends on

- a. High-resolution signal interpolation
- b. Compressing facial expressions
- c. Eye movement detection in grayscale images **×**
- d. Unique physical traits such as fingerprints or iris patterns

The correct answer is: Unique physical traits such as fingerprints or iris patterns

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:42:00	Saved: Eye movement detection in grayscale images	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

## Question 23

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q1. Reading an image in grayscale using OpenCV results in

- a. An image with inverted pixel values ✘
- b. A single matrix with 8-bit depth per pixel
- c. A multi-channel RGB image
- d. A thresholded binary image

The correct answer is: A single matrix with 8-bit depth per pixel

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:25:14	Saved: An image with inverted pixel values	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

Question **24**

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q16. HSV conversion is valuable because it

- a. Avoids resizing before filtering
- b. Applies inverse gamma correction to brightness ✘
- c. Enhances grayscale contrast
- d. Removes dependency on intensity values for color recognition

The correct answer is: Removes dependency on intensity values for color recognition

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:26:20	Saved: Applies inverse gamma correction to brightness	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

## Question 25

Correct

Mark 1.00 out of  
1.00

Hide sidebars

Q12. CNN deployment often requires model optimization such as

- a. Depth augmentation
- b. Histogram balancing
- c. Quantization and pruning ✓
- d. Recursive padding

The correct answer is: Quantization and pruning

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:27:00	Saved: Quantization and pruning	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Correct	1.00

## Question 26

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q13. Histogram of Oriented Gradients (HOG) feature extraction is useful for

- a. Texture synthesis ✖
- b. Capturing edge direction and magnitude for classification
- c. Thresholding image intensities
- d. Blurring large objects

The correct answer is: Capturing edge direction and magnitude for classification

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:27:42	Saved: Texture synthesis	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

## Question 27

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q11. Object recognition tasks generally involve

- a. Visualizing optical illusions using filters
- b. Counting edges per image block
- c. Enhancing images using global contrast ✘
- d. Identifying and localizing instances of known classes

The correct answer is: Identifying and localizing instances of known classes

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:28:32	Saved: Enhancing images using global contrast	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

## Question 28

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q14. Edge linking relies on gradient direction by

- a. Applying PCA to image borders ✘
- b. Propagating edge continuation along the gradient normal
- c. Thresholding only edge magnitude
- d. Moving along RGB histograms

The correct answer is: Propagating edge continuation along the gradient normal

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:29:37	Saved: Applying PCA to image borders	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

## Question 29

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q6. The `cv2.Canny()` function outputs

- a. A binary edge map
- b. A histogram of pixel differences
- c. A single RGB matrix
- d. Color contours of gradient direction ✖

The correct answer is: A binary edge map

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:30:03	Saved: Color contours of gradient direction	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

## Question 30

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q13. Normalizing pixel values between 0 and 1 improves

- a. Loss convergence to binary entropy
- b. Noise injection for regularization ✖
- c. Gradient flow and training stability
- d. Activation saturation in ReLU

The correct answer is: Gradient flow and training stability

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:32:05	Saved: Noise injection for regularization	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

## Question 31

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q13. Canny edge detection performs poorly in highly textured areas due to

- a. Limited dynamic range
- b. Chromatic dominance in grayscale conversion
- c. Insufficient scale normalization ✘
- d. Excessive strong edge triggers at multiple orientations

The correct answer is: Excessive strong edge triggers at multiple orientations

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:33:13	Saved: Insufficient scale normalization	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

## Question 32

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Q8. Activation maps in CNNs reflect

- a. Feature matching scores from the output layer
- b. Noise levels in the preprocessing stage
- c. Receptive field intensities from raw image blocks ✖
- d. Filter responses across spatial locations

The correct answer is: Filter responses across spatial locations

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:45	Started	Not yet answered	
2	6/05/25, 12:03:43	Saved: Receptive field intensities from raw image blocks	Answer saved	
3	6/05/25, 12:42:33	Attempt finished	Incorrect	0.00

◀ <https://www.youtube.com/...>

Jump to...

yolo ▶

☰ Table of contents

&gt; General

&gt; Topic 1

&gt; Topic 2

&gt; Topic 3

&gt; Topic 4

&gt; Topic 5

&gt; Topic 6

&gt; Topic 7

&gt; Topic 8

&gt; Topic 9

> **Topic 10**

> **Topic 11**

> **Topic 12**

> **Topic 13**

> **Topic 14**

▼ **Topic 15**

**Midterm Exam** 

> **Topic 16**

> **Topic 17**

> **Topic 18**

Hide sidebars

Copyright © PUeLC. 2024-2025 All rights reserved.



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

[Data retention summary](#)