



[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

2خ [202120386 خوله اشرف حسن محمد DS&AI](#)

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

**State** Finished

**Completed on** Tuesday, 6 May 2025, 12:47 PM

**Time taken** 35 mins 53 secs

**Marks** 22.00/32.00

**Grade** 20.63 out of 30.00 (68.75%)

## Question 1

Correct

Mark 1.00 out of  
1.00

Hide sidebars

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

- a. Thresholding image intensities
- b. Capturing edge direction and magnitude for classification ✓
- c. Texture synthesis
- 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, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:12:48	Saved: Capturing edge direction and magnitude for classification	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

## Question 2

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q1. Reading an image in grayscale using OpenCV results in

- a. A single matrix with 8-bit depth per pixel
- b. A multi-channel RGB image
- c. An image with inverted pixel values ✖
- 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, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:13:49	Saved: An image with inverted pixel values	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Incorrect	0.00

## Question 3

Correct

Mark 1.00 out of  
1.00

Hide sidebars

Q6. Pooling layers in CNNs are used to

- a. Replace convolution layers in shallow models
- b. Increase the number of filters at each level
- c. Reduce dimensionality while retaining key features ✓
- 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, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:14:12	Saved: Reduce dimensionality while retaining key features	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

## Question 4

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

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

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

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

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:15:40	Saved: It removes fully connected layers from detection heads	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Incorrect	0.00

## Question 5

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

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

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

The correct answer is: A binary edge map

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:16:17	Saved: A binary edge map	Answer saved	
3	6/05/25, 12:45:58	Saved: Color contours of gradient direction	Answer saved	
4	6/05/25, 12:47:56	Attempt finished	Incorrect	0.00

## Question 6

Correct

Mark 1.00 out of  
1.00

Hide sidebars

Q14. Filters in CNNs are learned using

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

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

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:16:44	Saved: Gradient descent and backpropagation over convolutional layers	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

## Question 7

Correct

Mark 1.00 out of  
1.00

Hide sidebars

Q3. Digital image representation stores pixel intensities

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

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

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:17:30	Saved: In discrete grid positions with intensity values	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

Question 8

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q12. Biometrics often depends on

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

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

Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:18:01	Saved: Unique physical traits such as fingerprints or iris patterns	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

## Question 9

Correct

Mark 1.00 out of  
1.00

Hide sidebars

Q14. Edge linking relies on gradient direction by

- a. Thresholding only edge magnitude
- b. Applying PCA to image borders
- c. Propagating edge continuation along the gradient normal ✓
- 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, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:18:53	Saved: Propagating edge continuation along the gradient normal	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

## Question 10

Correct

Mark 1.00 out of  
1.00

Hide sidebars

Q11. Object recognition tasks generally involve

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

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

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:19:04	Saved: Identifying and localizing instances of known classes	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

## Question 11

Correct

Mark 1.00 out of  
1.00

Hide sidebars

Q15. A batch size of 64 during training implies

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

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

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:19:54	Saved: Weights are updated every 64 gradient calculations	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

## Question 12

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. Filter responses across spatial locations
- d. Receptive field intensities from raw image blocks ✖

The correct answer is: Filter responses across spatial locations

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:21:34	Saved: Receptive field intensities from raw image blocks	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Incorrect	0.00

## Question 13

Correct

Mark 1.00 out of  
1.00

Hide sidebars

Q1. Edge detection in image processing aims to extract

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

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

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:23:51	Saved: Boundary-like features that encode shape	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

## Question 14

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q1. Describing interest points using feature vectors helps in

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

The correct answer is: Matching semantically similar regions across images

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:25:07	Saved: Matching semantically similar regions across images	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

Question **15**

Correct

Mark 1.00 out of  
1.00

Hide sidebars

Q16. HSV conversion is valuable because it

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

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

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:25:20	Saved: Removes dependency on intensity values for color recognition	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

## Question 16

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. Thresholding techniques in segmentation ✖
- c. Convolutional layers in deep networks
- 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, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:26:07	Saved: Thresholding techniques in segmentation	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Incorrect	0.00

## Question 17

Correct

Mark 1.00 out of 1.00

Hide sidebars

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

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

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

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:27:13	Saved: 75-dimensional dot product results per location	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

## Question 18

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q12. CNN deployment often requires model optimization such as

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

The correct answer is: Quantization and pruning

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:27:35	Saved: Quantization and pruning	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

## Question 19

Correct

Mark 1.00 out of  
1.00

Hide sidebars

Q5. Dilation in binary images leads to

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

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

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:44:17	Saved: Enlargement of white regions, closing small black holes	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

## Question 20

Incorrect

Mark 0.00 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, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:31:04	Saved: Match histogram ranges in smoothed images	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Incorrect	0.00

## Question 21

Correct

Mark 1.00 out of  
1.00

Hide sidebars

Q19. A fully connected Dense layer with ReLU activation introduces

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

The correct answer is: Nonlinear transformations after convolution

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:33:00	Saved: Nonlinear transformations after convolution	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

Question **22**

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

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

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

The correct answer is: float32

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:34:51	Saved: uint8	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Incorrect	0.00

Question **23**

Correct

Mark 1.00 out of  
1.00

Hide sidebars

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

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

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

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:35:10	Saved: Create a visual box around matched areas	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

Question **24**

Correct

Mark 1.00 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. Images are flattened for RNN input
- c. Image patches are square and color coded
- d. Channel depth is preserved for grayscale images ✓

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

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:36:21	Saved: Channel depth is preserved for grayscale images	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

Question **25**

Correct

Mark 1.00 out of  
1.00

Hide sidebars

Q5. Gaussian smoothing precedes derivative estimation to

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

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

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:36:34	Saved: Suppress high-frequency noise that misleads gradient computation	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

## Question 26

Correct

Mark 1.00 out of  
1.00

Hide sidebars

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

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

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

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:37:03	Saved: Excessive strong edge triggers at multiple orientations	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

## Question 27

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. Lack capability for multi-class classification ✖
- c. Rely on LBP rather than CNNs
- d. Only recognize shape templates, not identity

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

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:37:56	Saved: Lack capability for multi-class classification	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Incorrect	0.00

## Question 28

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q19. Image colorization tasks require

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

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

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:38:15	Saved: Filling grayscale regions with plausible chromatic values	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

## Question 29

Incorrect

Mark 0.00 out of  
1.00

Hide sidebars

Q14. Boundary detection failures typically arise when

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

The correct answer is: Local contrast is insufficient

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:38:50	Saved: High frequency filters are applied	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Incorrect	0.00

## Question 30

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q8. Otsu's thresholding automatically

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

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

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:39:16	Saved: Maximizes inter-class variance for threshold separation	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

## Question 31

Correct

Mark 1.00 out of  
1.00

Hide sidebars

Q13. Normalizing pixel values between 0 and 1 improves

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

The correct answer is: Gradient flow and training stability

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:40:03	Saved: Gradient flow and training stability	Answer saved	
3	6/05/25, 12:47:56	Attempt finished	Correct	1.00

## Question 32

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. Apply gradient filters for text-based input
- c. Convert images to binary maps directly
- d. Use spatial hierarchies to extract multi-scale features

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

## Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:12:03	Started	Not yet answered	
2	6/05/25, 12:43:36	Saved: Operate on raw pixel intensities without transformation	Answer saved	
3	6/05/25, 12:47:56	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](#)