2ا	202111335 ابراهيم احمد جابر الطيطي_DS&AI
Started on	Tuesday, 6 May 2025, 12:02 PM
State	Finished
Completed on	Tuesday, 6 May 2025, 12:39 PM
Time taken	36 mins 57 secs
Marks	29.00/32.00
Grade	**27.19** out of 30.00 (**90.63**%)

Question 1

Correct

Mark 1.00 out of 1.00

Hide sidebars

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

a. Images are flattened for RNN input

b. Channel depth is preserved for grayscale images ✓

c. Feature maps use dropout as regularization

d. Image patches are square and color coded

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

Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:05:07	Saved: Channel depth is preserved for grayscale images	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 2

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. Gradient clipping across pooling regions
- 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, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:04:39	Saved: Spatial invariance across filters	Answer saved	
3	6/05/25, 12:05:04	Saved: {$a}	Not yet answered	
4	6/05/25, 12:06:48	Saved: Nonlinear transformations after convolution	Answer saved	
5	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 3

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q11. Object recognition tasks generally involve

a. Identifying and localizing instances of known classes ✓

b. Visualizing optical illusions using filters

c. Counting edges per image block

d. Enhancing images using global contrast

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

Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:07:21	Saved: Identifying and localizing instances of known classes	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 4

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q5. Dilation in binary images leads to

a. Intensity normalization across object boundaries

b. Enlargement of white regions, closing small black holes ✓

c. Subdivision of connected components

d. Skewness correction of foreground objects

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

Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:07:23	Saved: Enlargement of white regions, closing small black holes	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 5

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q1. Describing interest points using feature vectors helps in

a. Reducing color intensity variation
b. Compressing image resolution for lossy storage
c. Matching semantically similar regions across images ✓
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, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:06:05	Saved: Tracking object deformation over time	Answer saved	
3	6/05/25, 12:06:43	Saved: {$a}	Not yet answered	
4	6/05/25, 12:07:40	Saved: Matching semantically similar regions across images	Answer saved	
5	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 6

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q3. Digital image representation stores pixel intensities

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

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

Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:08:14	Saved: In discrete grid positions with intensity values	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 7

Correct

Mark 1.00 out of 1.00

Hide sidebars

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

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

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

Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:08:33	Saved: Capturing edge direction and magnitude for classification	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 8

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. 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, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:08:51	Saved: Uniform pixel intensity values	Answer saved	
3	6/05/25, 12:09:04	Saved: {$a}	Not yet answered	
4	6/05/25, 12:09:33	Saved: Boundary-like features that encode shape	Answer saved	
5	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 9

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, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:10:29	Saved: Removes dependency on intensity values for color recognition	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 10

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q12. Biometrics often depends on

a. Eye movement detection in grayscale images
b. Compressing facial expressions
c. Unique physical traits such as fingerprints or iris patterns ✓
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, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:10:32	Saved: Unique physical traits such as fingerprints or iris patterns	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 11

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q10. Haar Cascades cannot identify specific object instances because they

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

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

Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:12:28	Saved: Only recognize shape templates, not identity	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 12

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q11. Threshold ratio in edge hysteresis is typically designed to

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

The correct answer is: Balance edge sensitivity and noise tolerance

Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:14:32	Saved: Balance edge sensitivity and noise tolerance	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 13

Correct

Mark 1.00 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, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:16:22	Saved: Excessive strong edge triggers at multiple orientations	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 14

Correct

Mark 1.00 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, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:17:50	Saved: Convolution is done once and reused across tasks	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 15

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q8. Otsu's thresholding automatically

a. Adjusts RGB channels based on luminance

b. Minimizes pixel intensity values globally

c. Maximizes inter-class variance for threshold separation ✓

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, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:17:57	Saved: Maximizes inter-class variance for threshold separation	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 16

Correct

Mark 1.00 out of 1.00

Hide sidebars

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

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

The correct answer is: Convolutional layers in deep networks

Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:19:46	Saved: Convolutional layers in deep networks	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 17

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q14. Filters in CNNs are learned using

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

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

Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:19:49	Saved: Gradient descent and backpropagation over convolutional layers	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 18

Correct

Mark 1.00 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. A thresholded binary image

d. An image with inverted pixel values

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

Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:19:05	Saved: A thresholded binary image	Answer saved	
3	6/05/25, 12:19:44	Saved: {$a}	Not yet answered	
4	6/05/25, 12:21:36	Saved: A single matrix with 8-bit depth per pixel	Answer saved	
5	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 19

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q6. Pooling layers in CNNs are used to

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

The correct answer is: Reduce dimensionality while retaining key features

Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:19:16	Saved: Reduce dimensionality while retaining key features	Answer saved	
3	6/05/25, 12:19:20	Saved: {$a}	Not yet answered	
4	6/05/25, 12:22:12	Saved: Reduce dimensionality while retaining key features	Answer saved	
5	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 20

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, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:22:56	Saved: Filling grayscale regions with plausible chromatic values	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 21

Incorrect

Mark 0.00 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, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:23:58	Saved: A 3D output with reduced spatial resolution	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Incorrect**	**0.00**

Question 22

Correct

Mark 1.00 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, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:24:01	Saved: Receptive field intensities from raw image blocks	Answer saved	
3	6/05/25, 12:24:24	Saved: Filter responses across spatial locations	Answer saved	
4	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 23

Correct

Mark 1.00 out of 1.00

Hide sidebars

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

a. Color contours of gradient direction
b. A single RGB matrix
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, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:25:19	Saved: A binary edge map	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 24

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q14. Edge linking relies on gradient direction by

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

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

Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:27:13	Saved: Propagating edge continuation along the gradient normal	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 25

Correct

Mark 1.00 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, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:32:41	Saved: Local contrast is insufficient	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 26

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q5. Gaussian smoothing precedes derivative estimation to

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

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

Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:28:54	Saved: Suppress high-frequency noise that misleads gradient computation	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 27

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Q12. CNN deployment often requires model optimization such as

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

The correct answer is: Quantization and pruning

Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:32:48	Saved: Recursive padding	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Incorrect**	**0.00**

Question 28

Correct

Mark 1.00 out of 1.00

Hide sidebars

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

a. Transform patches into binary

b. Create a visual box around matched areas ✓

c. Replace edge maps with histograms

d. Annotate noise in grayscale gradients

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

Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:30:23	Saved: Create a visual box around matched areas	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 29

Incorrect

Mark 0.00 out of 1.00

Hide sidebars

Q15. A batch size of 64 during training implies

a. Weights are updated every 64 gradient calculations
b. The model trains for only 64 iterations
c. Accuracy is computed after every 64 epochs ✕
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, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:32:59	Saved: Accuracy is computed after every 64 epochs	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Incorrect**	**0.00**

Question 30

Correct

Mark 1.00 out of 1.00

Hide sidebars

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

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

The correct answer is: float32

Response history

Step	Time	Action	State	Marks
1	6/05/25, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:32:01	Saved: float32	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 31

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q1. CNNs excel in image analysis because they

- a. Use spatial hierarchies to extract multi-scale features ✓
- b. Convert images to binary maps directly
- c. Operate on raw pixel intensities without transformation
- 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, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:33:55	Saved: Use spatial hierarchies to extract multi-scale features	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

Question 32

Correct

Mark 1.00 out of 1.00

Hide sidebars

Q13. Normalizing pixel values between 0 and 1 improves

a. Activation saturation in ReLU
b. Loss convergence to binary entropy
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, 12:02:19	Started	Not yet answered	
2	6/05/25, 12:33:45	Saved: Gradient flow and training stability	Answer saved	
3	**6/05/25, 12:39:16**	**Attempt finished**	**Correct**	**1.00**

◄ https://www.youtube.com/…

Jump to...

yolo ►

Table of contents