

USN

--	--	--	--	--	--	--	--	--	--

BCS714A

Seventh Semester B.E./B.Tech. Degree Examination, Dec.2025/Jan.2026

Deep Learning

Time: 3 hrs.

Max. Marks: 100

*Note: 1. Answer any FIVE full questions, choosing ONE full question from each module.**2. M : Marks , L: Bloom's level , C: Course outcomes.*

Module – 1			M	L	C
Q.1	a.	Explain with a suitable diagram, how according to Hubel and Wiesel's discovery a simple cell in the primary visual cortex of a cat fires at different rates depending upon the orientation of the line shown to the cat.	8	L2	CO1
	b.	Explain LeNet-5 hierarchical architecture and its use in automating the reading of postal service ZIP codes.	6	L2	CO1
	c.	Explain how the correctness of the sentence "The program impeccably translated the text" is determined using traditional learning and deep learning representations.	6	L2	CO1
OR					
Q.2	a.	Explain one-hot representation of words for the vocabulary = {The, bat, sat, on, a, cat} and describe the process of finding whether a word in the vocabulary is an animal or not for a simple NLP task.	8	L2	CO1
	b.	Explain AlexNet's hierarchical architecture and its use in face detection.	6	L2	CO1
	c.	Describe word-vector representation of words. Given the vector for words as follows: King = [-0.9, 1.9, 2.2] Man = [-1.1, 2.4, 3.0] Woman = [-3.2, 2.5, 2.6] Queen = [-3.0, 2.1, 1.7] Explain the process for finding that, vector for queen will be closest to the relation King – Man + Woman = Queen.	6	L2	CO1
Module – 2					
Q.3	a.	Illustrate with a suitable example, the need for regularization in deep learning.	8	L2	CO2
	b.	Write the mathematical formulation for the following parameter norm penalties: i) Limiting model capacity ii) L^2 parameter regularization iii) L^1 regularization For each of the above, give the expression for the norm penalty term $\Omega(\theta)$ and explain in each case, the effect of adding these penalties to the learning models. Also write the equations for the following giving explanation of the terms there-in. I. Closed form of Ridge regression II. Soft Thresholding rule for lasso slope solution.	8	L2	CO2

	c.	Explain the need for data augmentation and how it is effective in object recognition.	4	L2	CO2
OR					
Q.4	a.	Differentiate between Batch Gradient Descent and Minibatch Gradient Descent.	8	L2	CO2
	b.	Write the algorithm for stochastic gradient descent with momentum. Explain the core idea for introducing momentum and how it helps optimization.	8	L2	CO2
	c.	Describe the key mechanics of the Adam Optimization Algorithm.	4	L2	CO2
Module – 3					
Q.5	a.	Perform full convolution (flip the Kernel) for the image $f(x, y) = \begin{bmatrix} \textcircled{3} & 2 \\ 1 & 4 \end{bmatrix}$ and the Kernel $h(x, y) = \begin{bmatrix} \textcircled{7} & 6 \\ 5 & 8 \end{bmatrix}$ where the circled values 3 and 7 in the image and kernel indicate their (0, 0) positions.	9	L3	CO3
	b.	The convolution result for an image $f(x, y) = \begin{bmatrix} \textcircled{1} & 2 \\ 3 & 4 \end{bmatrix}$ and Kernel $h(x, y) = \begin{bmatrix} \textcircled{5} & 6 \\ 7 & 8 \end{bmatrix}$ is given as $G_{\text{shared}} = \begin{bmatrix} \textcircled{5} & 16 & 12 \\ 22 & 60 & 40 \\ 21 & 52 & 32 \end{bmatrix}$ Perform unshared convolution using the above information.	8	L3	CO3
	c.	In the CNN architecture, the output of the convolution operation is a feature map which is further given as input to an activation function followed by pooling. Given the result of convolution as $g(x, y) = \begin{bmatrix} 2 & 3 & -2 & -3 \\ 3 & 10 & 5 & 2 \\ 1 & 5 & 5 & 1 \end{bmatrix}$ Perform the activation function operation using Rectified Linear Unit (ReLU) followed by Max Pooling (2 × 2, stride 1)	3	L3	CO3
OR					
Q.6	a.	Compute the output of full cross-correlation (without kernel flipping) for an input image $f(x, y) = \begin{bmatrix} \textcircled{4} & 2 \\ 3 & 1 \end{bmatrix}$ and a Kernel $h(x, y) = \begin{bmatrix} \textcircled{8} & 6 \\ 7 & 5 \end{bmatrix}$ where the circled values 4 and 8 indicate their (0, 0) positions.	9	L3	CO3

	b. Compute the result of full tiled convolution (flip the kernels) for an input image $f(x, y) = \begin{bmatrix} 4 & 3 \\ 2 & 1 \end{bmatrix}$ and Kernel $h(x, y) = \begin{bmatrix} 7 & 5 \\ 6 & 8 \end{bmatrix}$ where the circled values 4 and 7 indicate their (0, 0) positions. Consider Tile_A = h(x, y) (flipped) and Tile_B = Tile_A + 10	8	L3	CO3
	c. The result of convolution + ReLU in the CNN workflow is given as shown below: $G(\text{CONV} + \text{ReLU}) = \begin{bmatrix} 4 & 17 & 21 & 18 \\ 9 & 32 & 43 & 30 \\ 2 & 10 & 16 & 8 \end{bmatrix}$ Perform the operations of Max Pooling (2×2, stride 1) followed by reshaping the output of Max Pooling into a 1D vector suitable to be given to the next stage of the CNN workflow.	3	L3	CO3

Module - 4

Q.7	a. Draw a diagram to show how unfolding a computational graph helps in understanding a Recurrent Neural Network (RNN). For the sequence of three words "I play tennis", show the step by step computation of the basic RNN processing of sequential data using forward propagation. Sequence of three words : I play tennis Where I = 1, play = 2, tennis = 3 Use, i) Input - hidden weight $u_x = 0.5$ ii) Hidden - hidden weight $w_h = 0.8$ iii) Bias $b = 0.1$ iv) Initial state $h_0 = 0$	10	L3	CO4
	b. Compare Bidirectional RNN with standard RNN. Compute the Backward pass ($\leftarrow$) for the input sequence : "I am thrilled" where I = 1, am = 1, thrilled = 2 Use, i) Initial backward state: $hB3 = 0$ ii) Backward Input-Hidden weight : $u_B = 0.6$ iii) Backward Hidden-Hidden weight : $W_B = 0.8$ iv) Backward Bias = $0.2 = \beta$	10	L3	CO4

OR

Q.8	a. Describe the roles of encoder and decoder in RNN encoder-decoder architectures. Assume that the final encoder state h_3 for the input sentence : I play cricket = 0.937. If the classification weights per class for the classes [learning, sports, others] is given by $w_y = [w_1, w_2, w_3] = [0.5, 1.5, -0.5]$ and classification biases per class are given as $b_y = [b_1, b_2, b_3] = [0.0, -0.9, 0.0]$, then decode the input sequence to the class "sports".	10	L3	CO4
	b. Compare Deep Recurrent Networks and Recursive Neural Networks. Represent the below provided input sentence: "I study Deep Learning" as a tree using Recursive Network concept.	10	L3	CO4

Module – 5

Q.9	a.	Explain the following common natural language preprocessing steps: i) Stop word removal ii) Stemming	10	L2	CO5
	b.	Explain how high dimensional word vectors are plotted to map them to two or three dimensions.	10	L2	CO5
OR					
Q.10	a.	Explain the following performance metrics of deep learning NLP models: i) The area under the ROC curve ii) The confusion matrix.	10	L2	CO5
	b.	Explain with a toy example, how the ROC Auc metric is calculated.	10	L2	CO5
