

Roll No.

ANNA UNIVERSITY (UNIVERSITY DEPARTMENTS)

B.E. / B.Tech. / B. Arch. (Full Time) - END SEMESTER EXAMINATIONS, NOV/DEC 2024

COMPUTER SCIENCE AND ENGINEERING

VII Semester

CS6005 Deep Learning Techniques
(Regulations – 2018 RUSA)

Max. Marks: 100

Time: 3 hrs

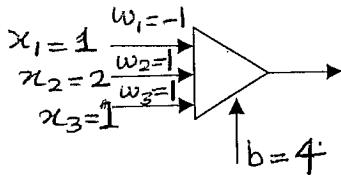
CO1	To learn the basic principles of supervised and unsupervised learning
CO2	To provide basic understanding of the concepts involved in deep learning
CO3	To understand the difference and similarities between the various forms of deep neural networks
CO4	To have knowledge about deep generative models
CO5	To know about the applications of deep learning techniques in various real-time problems

BL – Bloom's Taxonomy Levels

(L1 – Remembering, L2 – Understanding, L3 – Applying, L4 – Analysing, L5 – Evaluating, L6 – Creating)

PART- A (10 x 2 = 20 Marks)

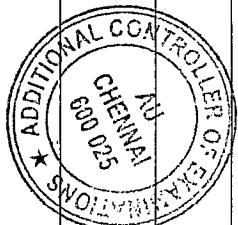
(Answer all Questions)

Q. No.	Questions	Marks	CO	BL
1.	When is a machine learning model said to be generalizable?	2	CO1	L2
2.	Suppose a neural network has an input layer, 2 hidden layers and one output layer. The input layer has 5 neurons, the hidden layers have 7 neurons each and the output layer has 3 neurons. Find the total number of parameters (weights and biases) that need to be trained in the neural network. Note that there are no biases in the input layer.	2	CO2	L4
3.	Consider the following neuron:  a) What is the cumulative input? b) What is the output if ReLU activation function is used?	2	CO2	L3
4.	What is multitask learning? When should it be used?	2	CO3	L4
5.	A single $(15 \times 15 \times 3)$ image is passed through a convolutional layer with 15 filters, each of size $(3 \times 3 \times 3)$. The padding size is 1 and the stride size is also 1. What is the size (that is, what are the dimensions) of the output?	2	CO3	L4
6.	What is meant by exploding gradient problem?	2	CO4	L2
7.	What is factor analysis?	2	CO5	L5
8.	Suppose you are building a model to diagnose COVID-19 using chest CT images. You are provided with 100,000 chest CT images, 1,000 of which are labelled. Which of the following learning techniques has the best chance of succeeding on this task? (i) Transfer learning or (ii) Supervised learning on the 1,000 labelled images? Why?	2	CO4	L2
9.	Why are restricted Boltzmann machines called as 'restricted'?	2	CO4	L2
10.	Suppose a video is to be given as input to a neural network. How should the video be represented?	2	CO5	L2

PART- B (8 x 8 = 64 Marks)

(Answer any 8 questions)

Q. No.	Questions	Marks	CO	BL																													
11.	a) Give the difference between supervised and unsupervised learning algorithms. Give one example for each category. b) Define Maximum Likelihood Estimation (MLE). Estimate the mean and variance of normal distribution using MLE.	2 6	CO1	L1																													
12.	a) Write in brief about the following activation functions: Leaky ReLU, Sigmoid b) Write in brief about the following loss functions: Mean Square Error Loss, KL Divergence	4 4	CO2	L2																													
13.	Explain the perceptron learning algorithm. Show how it can be used as a solution to the logic OR function.	8	CO2	L3																													
14.	Explain the working of Generative Adversarial Networks (GANs).	8	CO4	L4																													
15.	Explain the algorithm for the forward pass and the backward pass in a feedforward neural network.	8	CO3	L3																													
16.	With a neat diagram, explain the structure of an LSTM cell.	8	CO3	L3																													
17.	After training a neural network, you find that the training accuracy is very high (close to 100%) and the test accuracy is less (close to 50%). Which of the following methods would be more appropriate to reduce the difference? Why? Why would not the other methods be appropriate? (i) Dropout (ii) Generative Adversarial Networks (iii) Sigmoid activation function	8	CO3	L4																													
18.	a) Given an input image X of size (3,3) and a filter Y of size (2,2) as given below, what is the output after convolution X^*Y ? Assume stride = 1. X: <table border="1"><tr><td>1</td><td>2</td><td>0</td></tr><tr><td>0</td><td>1</td><td>2</td></tr><tr><td>2</td><td>2</td><td>0</td></tr></table> Y: <table border="1"><tr><td>1</td><td>-1</td></tr><tr><td>-2</td><td>1</td></tr></table> b) Find the output of 'average pooling' in the following matrix with a stride length of 2. <table border="1"><tr><td>1</td><td>0</td><td>3</td><td>10</td></tr><tr><td>0</td><td>1</td><td>1</td><td>6</td></tr><tr><td>1</td><td>2</td><td>4</td><td>5</td></tr><tr><td>3</td><td>1</td><td>3</td><td>2</td></tr></table> c) Given an input image of size 224×224 , a filter size of 5×5 and padding of 3, what are the possible values of stride?	1	2	0	0	1	2	2	2	0	1	-1	-2	1	1	0	3	10	0	1	1	6	1	2	4	5	3	1	3	2	3 3 2	CO3	L4
1	2	0																															
0	1	2																															
2	2	0																															
1	-1																																
-2	1																																
1	0	3	10																														
0	1	1	6																														
1	2	4	5																														
3	1	3	2																														
19.	Explain the architecture and working of any two types of RNNs. Give examples for each.	8	CO3	L3																													



20.	Write in brief about the following: a) Probabilistic PCA b) Independent Component Analysis	8	CO4	L2
21.	What is an autoencoder? Describe the approaches used in regularized autoencoders.	8	CO4	L4
22.	Explain the architecture and working of deep belief networks (DBN).	8	CO4	L4

PART- C (2 x 8 = 16 Marks)

(Answer all questions)

Q. No.	Questions	Marks	CO	BL
23.	You start training your model and notice underfitting, so you decide to add data augmentation as part of your preprocessing pipeline. Given that you are working with images of handwritten digits, for each data augmentation technique, state whether or not the following technique is appropriate for the task. If not, explain why not. (a) Scaling slightly (b) Flipping vertically (c) Flipping horizontally (d) Rotating by 90 or 180 degrees	8	CO4	L5
24.	Design a neural network that takes a lengthy video as input and generates an audio summary of the video in Tamil. Describe the input data set that you would feed to the neural network while training and testing.	8	CO5	L6

