


Name: Enrolment No:	
--------------------------------------	--

UPES
End Semester Examination, May 2023

Course: B. Tech. (CSE)
Program: Deep Learning (Neural Networks)
Course Code: CSAI3015P

Semester: VI
Time : 03 hrs.
Max. Marks: 100

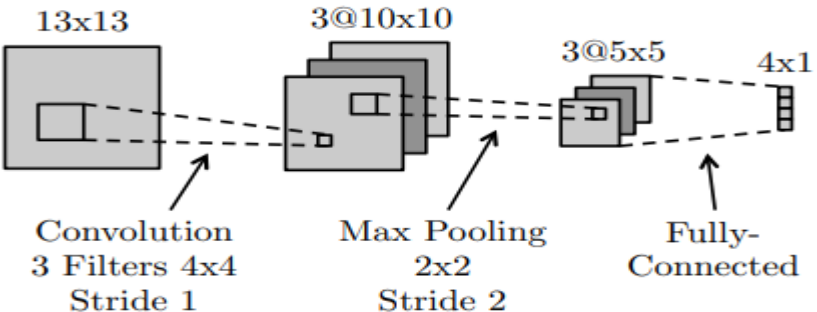
Instructions: Attempt all the questions.

SECTION A
(5Qx4M=20Marks)

S. No.	Question	Marks	CO
Q 1	Calculate the output y of a three-input neuron with bias. The input feature vector is $(x_1, x_2, x_3) = (0.9, 0.5, 0.3)$ and weight values are $[w_1, w_2, w_3, b] = [0.3, 0.2, -0.1, 0.4]$. Use binary Sigmoid function as activation function.	[4]	CO1
Q 2	Distinguish between overfitting and underfitting. How can it affect model generalization?	[4]	CO2
Q 3	Explain perceptron learning algorithm.	[4]	CO3
Q 4	Differentiate between supervised and unsupervised training. Explain with suitable examples.	[4]	CO1
Q 5	Do you think that a deep network is better than a shallow one? Justify your answer.	[4]	CO3

SECTION B
(4Qx10M= 40 Marks)

Q 6	Define confusion matrix. Suppose 10000 patients get tested for flu; out of them, 9000 are healthy and 1000 are sick. For the sick people, a test was positive for 620 and negative for 380. For the healthy people, the same test was positive for 180 and negative for 8820. Construct a confusion matrix for the data and compute the precision and recall for the data.	[10]	CO4															
Q 7	Can you represent the following boolean function with a single logistic threshold unit (i.e., a single unit from a neural network)? If yes, show the weights. If not, explain why not in 1-2 sentences. <table border="1" style="margin-left: auto; margin-right: auto; border-collapse: collapse;"> <thead> <tr> <th>A</th> <th>B</th> <th>f(A, B)</th> </tr> </thead> <tbody> <tr><td>1</td><td>1</td><td>0</td></tr> <tr><td>0</td><td>0</td><td>0</td></tr> <tr><td>1</td><td>0</td><td>1</td></tr> <tr><td>0</td><td>1</td><td>0</td></tr> </tbody> </table>	A	B	f(A, B)	1	1	0	0	0	0	1	0	1	0	1	0	[10]	CO2
A	B	f(A, B)																
1	1	0																
0	0	0																
1	0	1																
0	1	0																

Q 8	<p>What is generative deep learning? How does generative deep learning differ from discriminative deep learning? What are some common applications of generative deep learning?</p> <p style="text-align: center;">OR</p> <p>How does a long short-term memory (LSTM) cell work, and what are some advantages of using LSTMs over other types of RNNs?</p>	[10]	CO3
Q 9	<p>What is ensemble learning? Discuss the difference between bagging and boosting, and hard and soft voting in ensemble techniques?</p>		CO4
<p>SECTION-C (2Qx20M=40 Marks)</p>			
Q 10	<p>Explain the principle of the gradient descent algorithm. You are given the function $f(x, y) = x^2 + xy + y^2$ and are trying to find a local minimum using gradient descent. You randomly start with $x = 1.3$ and $y = 5.4$. Perform the first step of gradient descent with learning rate $= 0.01$. Show the resulting values for x and y as well as all of your calculations.</p>	[20]	CO4
Q11	<p>Below is a diagram of a small convolutional neural network that converts a 13x13 image into 4 output values. The network has the following layers/operations from input to output: convolution with 3 filters, max pooling, ReLu, and finally a fully connected layer. For this network we will not be using any bias/offset parameters (b). Please answer the following questions about this network.</p> <div style="text-align: center;">  <p style="text-align: center;"> 13x13 3@10x10 3@5x5 4x1 Convolution Max Pooling Fully-Connected 3 Filters 4x4 2x2 4x1 Stride 1 Stride 2 </p> </div> <p>a) How many weights in the convolutional layer do we need to learn? b) How many ReLu operations are performed on the forward pass? c) How many weights do we need to learn for the entire network?</p> <p style="text-align: center;">OR</p> <p>Answer the following questions.</p> <p>a) How does splitting a dataset into train, dev and test sets help identify overfitting? b) What Is the Cost Function? c) What Are Hyperparameters? d) What Is the Difference Between Batch Gradient Descent and Stochastic Gradient Descent?</p>	[20]	CO2