



Roll No:

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

BTECH
(SEM VII) THEORY EXAMINATION 2024-25
DEEP LEARNING

TIME: 3 HRS

M.MARKS: 100

Note: Attempt all Sections. In case of any missing data; choose suitably.

SECTION A

1. Attempt all questions in brief. 2 x 10 = 20

Q no.	Question	CO	Level
a.	What is a loss function, and why is it essential for training a neural network?	1	K1
b.	What is stochastic gradient descent?	1	K4
c.	Explain the significance of the "AI Winter" in deep learning history.	2	K2
d.	What is the probabilistic theory of deep learning?	2	K1
e.	Define Linear Discriminant Analysis (LDA).	3	K1
f.	Explain how distance metrics are used in machine learning.	3	K2
g.	Define non-convex optimization.	4	K1
h.	What is generalization in deep learning?	4	K1
i.	What is WaveNet, and why is it significant for audio generation?	5	K1
j.	What is Word2Vec?	5	K1

SECTION B

2. Attempt any three of the following: 10 x 3 = 30

a.	Illustrate how activation functions like ReLU and Sigmoid work in a shallow network.	1	K3
b.	Explain the role of the generator and discriminator in GANs.	2	K2
c.	What is the significance of the AlexNet architecture in deep learning history?	3	K1
d.	Explain the role of STNs in addressing spatial invariance in deep learning.	4	K2
e.	Describe the role of deep learning in bioinformatics applications.	5	K2

SECTION C

3. Attempt any one part of the following: 10 x 1 = 10

a.	Explain the purpose of backpropagation in training neural networks.	1	K2
b.	Describe the role of stochastic gradient descent (SGD) in optimization.	1	K2

4. Attempt any one part of the following: 10 x 1 = 10

a.	Demonstrate the working of a CNN on image classification tasks.	2	K3
b.	Demonstrate how batch normalization reduces internal covariate shift.	2	K3

5. Attempt any one part of the following: 10 x 1 = 10

a.	Demonstrate the effect of Xavier and He initialization on training a ConvNet.	3	K3
b.	Illustrate the role of residual connections in ResNet and their importance.	3	K3

6. Attempt any one part of the following: 10 x 1 = 10

a.	Explain how LSTMs address the vanishing gradient problem in RNNs.	4	K2
b.	Explain how RNNs are used for next-word prediction in text data.	4	K2

7. Attempt any one part of the following: 10 x 1 = 10

a.	Describe how attention mechanisms improve the performance of image captioning models.	5	K2
b.	Explain how semantic segmentation helps in scene understanding.	5	K2