R22 Regulation *Subject code:4E7EA*

TKR COLLEGE OF ENGINEERING AND TECHNOLOGY

(Autonomous, Accredited by NAAC with 'A+' Grade)

B.Tech VII Semester Regular/Supplementary Examinations, October 2026

DEEP LEARNING

(CSE)

Maximum Marks: 60 Date: 26.10.2026 Duration: 3 hours

Note: **1.This question paper contains two parts A and B.**
2. Part A is compulsory which carries 10 marks. Answer all questions in Part A.
3. Part B consists of 5 Units. Answer any one full question from each unit.
4. Each question carries 10 marks and may have a, b, c, d as sub questions.

Part-A

All the following questions carry equal marks (10X1M=10 Marks)		Marks	CO	BTL
1.a)	A machine learning model performs very well on the training data but gives poor results on unseen validation data. Identify the learning problem observed in this situation.	1M	1	L1
b)	In a deep feed-forward network, what is the primary purpose of back-propagation?	1M	1	L2
c)	A neural network is trained with a very large number of parameters and begins to memorize the training data. Name one regularization technique that can be used to reduce this problem.	1M	2	L1
d)	What is the purpose of early stopping while training a deep neural network?	1M	2	L2
e)	In a convolutional neural network, which operation is commonly used to reduce the spatial dimensions of feature maps?	1M	3	L1
f)	State one important advantage of using convolution operations for processing image data.	1M	3	L2
g)	Why is an ordinary RNN likely to face difficulty in learning long-term dependencies in sequential data?	1M	4	L2
h)	Name the gated recurrent architecture specifically designed to address the vanishing-gradient problem and capture long-term dependencies.	1M	4	L1
i)	In evaluating a deep learning model, why should a performance metric be selected according to the application rather than relying only on accuracy?	1M	5	L2
j)	Mention one practical strategy that can be considered when a deep learning model does not provide satisfactory performance because of insufficient training data.	1M	5	L2

Part-B				
Answer All the following questions. (5X10M=50Marks)		Marks	CO	BTL
2	a) Explain the concepts of overfitting, underfitting, bias, variance, hyperparameters, and validation sets in the context of developing the proposed machine learning model. b) Demonstrate how stochastic gradient descent and back-propagation can be used to train a deep feed-forward neural network. Explain the role of hidden units and network architecture during the learning process.	5M 5M	1	L2, L3
	OR			
3	a) Explain why a single-layer neural network cannot effectively learn the XOR problem and describe how hidden units and multiple layers overcome this limitation. b) Apply the concept of gradient-based learning and back-propagation to describe how the weights of a multi-layer feed-forward network can be updated during training.	5M 5M	1	L2, L3
4	Explain how parameter norm penalties, dataset augmentation, early stopping, and dropout help in controlling overfitting in deep neural networks.	10M	2	L2
	OR			
5	a) Describe the major challenges in neural network optimization and explain the importance of parameter initialization in training deep models. b) Explain the adaptive learning-rate strategies to improve the training process.	5M 5M	2	L2, L3
6	a) Explain the convolution operation and pooling operation and describe why they are useful for extracting features from image data. b) Demonstrate how convolution and pooling can be applied sequentially to an image to obtain progressively useful feature representations.	5M 5M	3	L2, L3
	OR			
7	a) Explain the different variants of the basic convolution function and describe the role of convolutional networks in handling structured image data. b) Apply the concepts of convolution, pooling, and efficient convolution algorithms to design a suitable feature-extraction pipeline for the product-inspection system.	5M 5M	3	L2, L3
8	Explain the working of a Recurrent Neural Network (RNN) and describe how an RNN processes sequential information using an unfolded computational graph.	10M	4	L3
	OR			
9	a) Explain the architecture and working of an encoder-decoder sequence-to-sequence model. Describe the roles of the encoder and decoder in processing sequential data.	5M	4	L2, L3

	b) Apply the concept of LSTM to the translation problem and demonstrate how its gated structure helps in retaining important information over long sequences.	5M		
10	Explain the importance of performance metrics, default baseline models, hyperparameter selection, and debugging strategies in developing a reliable deep learning system.	10M	5	L2
	OR			
11	Describe the practical methodology involved in developing a multi-digit number recognition system, including model selection, performance evaluation, hyperparameter selection, and debugging and apply deep learning methodology to the proposed multi-digit recognition problem and explain how the same general approach can be adapted for computer vision, speech recognition, and natural language processing applications.	10M	5	L3