23CSPC401

USN | | | | | | | | | | |

MANGALORE INSTITUTE OF TECHNOLOGY & ENGINEERING

(A Unit of Rajalaxmi Education Trust®, Mangalore)
Autonomous Institute affiliated to VTU, Belagavi, Approved by AICTE, New Delhi
Accredited by NAAC with A+ Grade & ISO 9001:2015 Certified Institution

Model Question Paper

Seventh Semester B.E Degree Examination

Deep Learning

Time: 3 Hours (180 Minutes) **Max. Marks: 100**

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

		Module -1	M	L	C
Q1	a.	What is Deep Learning? Give the difference between Machine Learning and Deep Learning.	6	L2	CO1
	b.	Discuss the setup and initialization issues in training deep neural networks. Explain why proper weight initialization is important for effective learning.	8	L2	CO1
	c.	Explain different gradient descent strategies used in training deep neural networks.	6	L2	CO1
		OR			
Q2	a.	Explain the backpropagation algorithm used for training deep neural networks, including the role of the chain rule in computing gradients.	6	L2	CO1
	b.	Explain the Vanishing Gradient Problem and Exploding Gradient Problem. What causes these problems and how do they affect the training of deep neural networks?	8	L2	CO1
	c.	Explain the working principle of Batch Normalization and discuss how it improves training speed and stability in deep neural networks.	6	L2	CO1
		Module- 2			
Q3	a.	A deep neural network achieves 99% training accuracy but only 68% test accuracy. Analyze this scenario and apply the bias-variance concept to identify the problem and suggest two possible remedies.	10	L3	CO2
	b.	A dataset has only 500 labeled samples, and the team wants to tune hyperparameters reliably without wasting data. Apply an appropriate model evaluation strategy discussed in the text and justify your choice.	10	L3	CO2
		OR			
Q4	a.	You are training a neural network and observe that the L1-regularized model produces many zero-valued weights, while the L2-regularized model produces small but non-zero weights. Apply this observation to decide which regularization method you would use if you wanted an interpretable, sparse network, and justify your answer.	10	L3	CO2
	b.	During training, the validation loss decreases steadily up to epoch 40 and then starts increasing, while training loss keeps decreasing. Apply the early stopping technique to this scenario and explain how you would decide the stopping point and why it improves generalization.	10	L3	CO2
		Module – 3			
Q5	a.	Consider a 1-dimensional time-series with values 2, 1, 3, 4, 7. Perform a convolution with a 1-dimensional filter 1, 0, 1 and zero padding.	10	L3	CO3
	b.	Illustrate the convolution layer and pooling layer in a CNN with suitable diagrams.	10	L3	CO3
		OR			

Q6	a.	Compare VGG and ResNet architectures and apply this comparison to justify why ResNet can be trained with a much greater depth (e.g., 152 layers) than VGG (19 layers) without degrading performance.	10	L3	CO3
	b.	Design a Convolutional Neural Network (CNN) for a multi-class image classification problem using convolutional layers, ReLU activation, max-pooling, and a fully connected output layer. Justify the choice of each component in the architecture.	10	L3	CO3
Module – 4					
Q7	a.	Illustrate the architecture of Recurrent Neural Networks.	10	L3	CO4
	b.	Apply the concepts of Recurrent Neural Networks (RNNs) to design a Question-Answering System and justify why an RNN-based architecture is suitable for this application.	10	L3	CO4
OR					
Q8	a.	A standard RNN is trained on long text sequences (500+ words) for language modeling, but it fails to capture dependencies between words that are far apart in the sequence. Apply your understanding of RNN training challenges to identify the problem and propose a solution.	10	L3	CO4
	b.	Given a computational budget constraint where training time must be minimized without significantly sacrificing accuracy, apply your knowledge of gated recurrent architectures to choose between LSTM and GRU, and justify your choice.	10	L3	CO4
Module – 5					
Q9	a.	Illustrate the basic framework of Reinforcement Learning including agent, environment, and reward.	6	L3	CO4
	b.	In a game-playing agent, you want to update the value of taking an action without waiting for the entire game to finish. Apply the concept of bootstrapping to design an update rule the agent could use after each single move.	10	L3	CO4
	c.	A robot must move from a start cell to a goal cell in a grid. Apply the basic RL framework to identify its state, action, and reward	4	L3	CO4
OR					
Q10	a.	Illustrate the role of Monte Carlo Tree Search in decision-making problems.	6	L3	CO4
	b.	A self-driving car agent receives raw camera pixel images as its state input, making the state space extremely high-dimensional. Apply the concept of Deep Q-Networks to explain how the agent could learn an effective driving policy from this input.	10	L3	CO4
	c.	A dog is being trained using treats: it gets a treat when it sits on command, and no treat when it doesn't. Apply the concept of reinforcement learning to explain what represents the agent, environment, action, and reward in this scenario.	4	L3	CO4

-*****-