



SRM VALLIAMMAI ENGINEERING COLLEGE

SRM Nagar, Kattankulathur – 603 203.

(An Autonomous Institution)

DEPARTMENT OF COMPUTER APPLICATIONS

QUESTION BANK



III SEMESTER

PMC102 NEURAL NETWORK AND DEEP LEARNING

Regulation – 2024

Academic Year 2026- 2027

Prepared by

Ms.G.S.Gayathri, Assistant Professor/MCA

SRM VALLIAMMAI ENGINEERING COLLEGE

DEPARTMENT OF COMPUTER APPLICATIONS

QUESTION BANK

SUBJECT : PMC102-NEURAL NETWORK AND DEEP LEARNING

SEM / YEAR: III / II

UNIT I INTRODUCTION				
Biological Neuron, Idea of computational units, McCulloch–Pitts unit and Thresholding logic, Linear Perceptron, Perceptron Learning Algorithm, Linear separability, Convergence theorem for Perceptron Learning Algorithm, Type of network architecture, Activation functions, Basic Learning rules.				
Q.No.	Question	Bloom's Taxonomy Level	Competence	Course Outcome
PART – A				
1.	Define a biological neuron and name its main components.	BT1	Remembering	CO 1
2.	What is the learning rate in the Perceptron Learning Algorithm? Why is it important?	BT1	Remembering	CO 1
3.	Define a computational unit (artificial neuron) in Artificial Neural Networks?	BT1	Remembering	CO 1
4.	Explain the analogy between a biological neuron and an artificial neuron.	BT2	Understanding	CO 1
5.	Define a synaptic weight and state what it represents in an artificial neuron.	BT1	Remembering	CO 1
6.	Describe the McCulloch–Pitts (M-P) neuron model.	BT2	Understanding	CO 1
7.	List the assumptions made in the McCulloch–Pitts neuron model.	BT1	Remembering	CO 1
8.	Explain thresholding logic and its significance in neural computation.	BT2	Understanding	CO 1
9.	Demonstrate how a McCulloch–Pitts neuron can implement AND and OR logic gates.	BT2	Understanding	CO 1
10.	Find the output of an M-P neuron for $x_1 = 1, x_2 = 1, w_1 = 1, w_2 = 1, \theta = 2$.	BT1	Remembering	CO 1
11.	Define a perceptron and state its basic purpose.	BT1	Remembering	CO 1
12.	Write the mathematical equation of a linear perceptron output.	BT1	Remembering	CO 1
13.	Explain the role of the bias term in a perceptron.	BT2	Understanding	CO 1
14.	Compare a McCulloch–Pitts neuron with a perceptron.	BT2	Understanding	CO 1
15.	Outline the Perceptron Learning Algorithm and its weight-update rule.	BT2	Understanding	CO 1
16.	Why is the learning rate important in the Perceptron Learning Algorithm?	BT1	Remembering	CO 1
17.	Given $w = [0.2 \ 0.3]$, $x = [1 \ 1]$, $t = 1$, $y = 0$, and $\eta = 0.1$. Compute the updated weights using the perceptron learning rule.	BT2	Understanding	CO 1
18.	Illustrate the concept of linear separability with an example.	BT2	Understanding	CO 1

19.	Explain why a single-layer perceptron cannot solve the XOR problem.	BT2	Understanding	CO 1
20.	State the Perceptron Convergence Theorem and the condition required for convergence.	BT1	Remembering	CO 1
21.	List different types of neural network architectures and name any two.	BT1	Remembering	CO 1
22.	Contrast feedforward and feedback (recurrent) neural networks.	BT2	Understanding	CO 1
23.	Differentiate between single-layer and multi-layer neural networks.	BT2	Understanding	CO 1
24.	List four commonly used activation functions and explain the need for non-linear functions.	BT1 + BT2	Remembering + Understanding	CO 1
25.	Find the output for $v = 2$ using a binary step activation function with $\theta = 1$, and state the Hebbian learning rule.	BT1	Remembering	CO 1
PART – B				
1.	(a) Examine the structure and functioning of a biological neuron with a neat diagram. (8 Marks) (b) Compare a biological neuron with an artificial neuron by mapping their corresponding components. (8 Marks)	BTL-4	Analyzing	CO 1
2.	Construct a computational model of an artificial neuron by incorporating inputs, weights, bias, summation function, and activation function. Demonstrate its working with a suitable example.	BTL-3	Applying	CO 1
3.	(a) Explain the McCulloch–Pitts (M-P) neuron model with its mathematical formulation, assumptions, and thresholding logic. (8 Marks) (b) Design M-P neurons to implement AND, OR, and NOT logic gates. (8 Marks)	BTL-6	Creating	CO 1
4.	(a) Analyze the working of thresholding logic with suitable examples. (8 Marks) (b) Examine why the McCulloch–Pitts neuron cannot implement the XOR function. (8 Marks)	BTL-4	Analyzing	CO 1
5.	Construct the architecture and mathematical model of a Linear Perceptron. Formulate its output equation and explain the roles of weights, bias, and activation function in binary classification.	BTL-3	Applying	CO 1
6.	Develop a flowchart for the Perceptron Learning Algorithm. Formulate the weight-update rule and demonstrate the learning procedure using a suitable numerical example.	BTL-3	Applying	CO 1
7.	Apply the Perceptron Learning Algorithm to the given initial weights, bias, learning rate, and training patterns, and compute the updated weights after each iteration until convergence.	BTL-3	Applying	CO 1
8.	Evaluate the conditions necessary for the Perceptron Convergence Theorem and discuss its significance in neural-network learning.	BTL-5	Evaluating	CO 1
9.	Explain the concept of linear separability with suitable examples and diagrams. Analyze why the XOR problem is	BTL-4	Analyzing	CO 1

	not linearly separable and explain how a multilayer perceptron overcomes this limitation.			
10.	Construct neat diagrams representing single-layer, multilayer, feedforward, and feedback/recurrent neural-network architectures. Compare their characteristics and applications.	BTL-4	Analyzing	CO 1
11.	Contrast feedforward and recurrent neural-network architectures in terms of structure, information flow, memory, learning capability, and applications.	BTL-4	Analyzing	CO 1
12.	(a) Explain the concept of an activation function and state its significance in an artificial neuron. (8 Marks) (b) Develop representations of Binary Step, Sigmoid, Tanh, ReLU, Leaky ReLU, Parametric ReLU, and Softmax, including mathematical expressions, graphs, characteristics, and applications. (8 Marks)	BTL-3	Applying	CO 1
13.	Evaluate the suitability of Binary Step, Sigmoid, Tanh, ReLU, Leaky ReLU, Parametric ReLU, and Softmax activation functions for different neural-network applications. Justify the selections with appropriate reasoning.	BTL-5	Evaluating	CO 1
14.	Evaluate the suitability of Hebbian, Perceptron, Delta, Correlation, and Competitive Learning rules for different neural-network applications, with their mathematical formulations and suitable examples.	BTL-6	Creating	CO 1
15.	(a) Evaluate the suitability of Binary Step, Sigmoid, Tanh, ReLU, Leaky ReLU, Parameteric ReLU and Softmax activation functions for different neural network applications. Justify your choice with appropriate reasoning. (8 Marks) (b) Compare the Sigmoid and Softmax activation functions using a suitable example. (8 Marks)	BTL-4	Analyzing	CO 1
16.	Design a perceptron-based classifier for a bank loan approval system using customer income, credit score, and repayment history. Develop its learning process and evaluate its limitations for non-linearly separable data.	BTL-6	Creating	CO 1
17.	Invent a modified computational neuron capable of processing real-valued inputs by extending the McCulloch–Pitts model. Justify its design, explain its working, and compare it with the traditional M-P neuron.	BTL-6	Creating	CO 1
18.	Evaluate the evolution of computational neuron models from the biological neuron to the McCulloch–Pitts neuron and perceptron. Assess the importance of the Perceptron Convergence Theorem in modern neural-network theory.	BTL-5	Evaluating	CO 1


UNIT II FEEDFORWARD NETWORKS


Multilayer Neural Network, Gradient Descent learning, Back propagation, Empirical Risk Minimization, regularization, Radial Basis Neural Network

Q.No.	Question	Bloom's Taxonomy Level	Competence	Course Outcome
-------	----------	------------------------	------------	----------------

PART – A

1.	State the meaning of a Multilayer Neural Network.	BTL-1	Remembering	CO 2
2.	Sketch the architecture of a Multilayer Feedforward Neural Network and label its major components.	BTL-2	Understanding	CO 2
3.	Identify the different layers present in a Multilayer Neural Network.	BTL-1	Remembering	CO 2
4.	Summarize the advantages of a Multilayer Neural Network over a Single-layer Perceptron.	BTL-2	Understanding	CO 2
5.	State the meaning of Gradient Descent Learning.	BTL-1	Remembering	CO 2
6.	Recall the weight-update equation used in Gradient Descent Learning.	BTL-1	Remembering	CO 2
7.	State the purpose of the learning rate (η) in Gradient Descent.	BTL-2	Understanding	CO 2
8.	Distinguish between Batch Gradient Descent and Stochastic Gradient Descent.	BTL-2	Understanding	CO 2
9.	State the meaning of the Backpropagation Algorithm.	BTL-1	Remembering	CO 2
10.	Enumerate the major steps involved in the Backpropagation Algorithm.	BTL-1	Remembering	CO 2
11.	Explain why Backpropagation is used in Multilayer Neural Networks.	BTL-2	Understanding	CO 2
12.	Describe the role of the activation function in the Backpropagation Algorithm.	BTL-2	Understanding	CO 2
13.	What is overfitting? How does regularization help prevent it?	BTL-1	Remembering	CO 2
14.	Explain the objective of Empirical Risk Minimization in machine learning.	BTL-2	Understanding	CO 2
15.	State the meaning of Regularization in neural networks.	BTL-1	Remembering	CO 2
16.	Why is regularization used in neural networks?	BTL-2	Understanding	CO 2
17.	Distinguish between L1 and L2 Regularization.	BTL-2	Understanding	CO 2
18.	Using the Gradient Descent update rule, compute the updated weight if the old weight is 0.6, the learning rate is 0.1, and the gradient is 0.4.	BTL-2	Understanding	CO 2
19.	State the meaning of a Radial Basis Function (RBF) Neural Network.	BTL-1	Remembering	CO 2
20.	Sketch the architecture of a Radial Basis Function Neural Network and label its layers.	BTL-2	Understanding	CO 2
21.	Identify the main components of a Radial Basis Function Neural Network.	BTL-1	Remembering	CO 2
22.	Given the error ($E = 0.5(t-o)^2$), target output ($t=1$), and actual output ($o=0.8$), calculate the error.	BTL-2	Understanding	CO 2
23.	Calculate the error for the given values using $E = \frac{1}{2}(t - o)^2$, where target output $t = 1$ and actual output $o = 0.8$.	BTL-2	Understanding	CO 2
24.	Given the error function $E = 0.5(t - o)^2$, target output $t = 0.9$, and actual output $o = 0.6$, calculate the error.	BTL-2	Understanding	CO 2
25.	An RBF neuron has center $c = 2$, input $x = 3$, and width $\sigma = 1$. Write the Gaussian activation function and determine the exponent value.	BTL-1	Remembering	CO 2

PART –B				
1.	Illustrate the architecture and working of a Multilayer Neural Network (MLNN) with a neat diagram. Explain how information propagates from the input layer to the output layer.	BTL-3	Applying	CO 2
2.	Analyze the architecture of a Multilayer Perceptron (MLP). Demonstrate how hidden layers improve the learning capability of neural networks with suitable examples such as XOR classification.	BTL-4	Analyzing	CO 2
3.	Consider the following Multilayer Perceptron (MLP) with 2 input neurons, 2 hidden neurons, and 2 output neurons. The initial input values, weights, and biases are given in the diagram. Using the given network, calculate the final output values o_1 and o_2 separately by using Sigmoid, ReLU and Tanh activation function. Show the complete forward propagation calculations for each activation function. a) , b) 	BTL-5	Evaluating	CO 2
4.	Derive the Gradient Descent Learning Algorithm. Illustrate the role of the learning rate and analyze the factors affecting convergence.	BTL-4	Analyzing	CO 2
5.	(a) Analyze the Gradient Descent Learning Algorithm and illustrate the parameter-update process. (8 Marks) (b) Apply the Backpropagation Algorithm and derive the weight-update equations for the hidden and output layers. (8 Marks)	BTL-4	Analyzing	CO 2
6.	Explain the different types of neural network architectures (single-layer, multilayer, feedforward, and feedback/recurrent) with neat diagrams. Compare their characteristics and applications.	BTL-3	Applying	CO 2

7.	Evaluate the importance of Backpropagation in training multilayer neural networks. Discuss its advantages, limitations, and convergence issues.	BTL-5	Evaluating	CO 2
8.	Analyze the concept of Empirical Risk Minimization (ERM) and examine its significance in supervised learning. Discuss its limitations with suitable examples.	BTL-4	Analyzing	CO 2
9.	Apply the Backpropagation Algorithm for the given network architecture, initial weights, learning rate, and training sample. Compute the updated weights after one iteration. a) , b) 	BTL-3	Applying	CO 2
10.	Analyze the concept of Regularization in neural networks. Differentiate L1 and L2 Regularization based on their objective functions, effect on weights, sparsity, and model performance.	BTL-4	Analyzing	CO 2
11.	Evaluate L1, L2, Early Stopping, and Dropout as techniques for reducing overfitting. Recommend suitable techniques for different neural-network applications with justification.	BTL-5	Evaluating	CO 2
12.	Analyze the architecture and working principle of a Radial Basis Function (RBF) Neural Network. Compare its learning mechanism with that of an MLP.	BTL-4	Analyzing	CO 2
13.	Apply the Gaussian radial basis function to the given center, width, and input vector. Calculate the activation value and determine the network output using the given weights.	BTL-3	Applying	CO 2
14.	Evaluate whether an MLP or RBF Network is more suitable for a hospital-based disease diagnosis system. Justify your recommendation based on training, computational complexity, nonlinear classification, and generalization.	BTL-5	Evaluating	CO 2
15.	Evaluate the causes of overfitting in an image-recognition system. Select and justify suitable regularization techniques to improve its generalization performance.	BTL-5	Evaluating	CO 2
16.	a) Explain the Gradient Descent Learning Algorithm with	BTL-4	Analyzing /	CO 2

	suitable illustrations. b) Compare Batch Gradient Descent, Stochastic Gradient Descent (SGD), and Mini-batch Gradient Descent in terms of convergence speed, computational complexity, and applications.	/ BTL-3	Applying	
17.	(a) Analyze Empirical Risk Minimization and examine its relationship with training error and generalization error. (8 Marks) (b) Evaluate different Regularization Techniques used to reduce overfitting, including L1, L2, Early Stopping, and Dropout. (8 Marks)	BTL-4 / BTL-5	Analyzing / Evaluating	CO 2
18.	Design a Feedforward Neural Network for handwritten digit recognition. Select and justify an appropriate architecture, activation functions, learning algorithm, loss function, and regularization technique.	BTL-6	Creating	CO 2

UNIT III RECURRENT NEURAL NETWORKS

Back propagation through time, Long Short Term Memory, Gated Recurrent Units, Bidirectional LSTMs and Bidirectional RNNs.

Q.No.	Question	Bloom's Taxonomy Level	Competence	Course Outcome
PART – A				
1	State the advantages of Recurrent Neural Networks over Feedforward Neural Networks.	BTL-1	Remembering	CO 3
2	List the major steps involved in the BPTT algorithm.	BTL-1	Remembering	CO 3
3	Identify the two major gradient problems encountered during BPTT.	BTL-1	Remembering	CO 3
4	Why BPTT becomes computationally expensive for very long sequences.	BTL-2	Understanding	CO 3
5	Explain how gradient clipping helps overcome the exploding gradient problem.	BTL-2	Understanding	CO 3
6	State the purpose of the cell state in an LSTM network.	BTL-1	Remembering	CO 3
7	Identify the three gates of a standard LSTM cell.	BTL-1	Remembering	CO 3
8	Describe how the Forget Gate controls information retained in the LSTM cell state.	BTL-2	Understanding	CO 3
9	Explain the role of the Input Gate in an LSTM.	BTL-2	Understanding	CO 3
10	Interpret the effect of a Forget Gate output of 0 on the previous cell-state information.	BTL-2	Understanding	CO 3
11	Numerical: If $f_t = 0.8$ and $C_{t-1} = 0.5$, calculate the retained contribution $f_t C_{t-1}$.	BTL-2	Understanding	CO 3
12	State the two gates used in a Gated Recurrent Unit (GRU).	BTL-1	Remembering	CO 3
13	Identify one major architectural difference between a GRU and an LSTM.	BTL-1	Remembering	CO 3
14	Explain how the Update Gate in a GRU helps retain information from previous time steps.	BTL-2	Understanding	CO 3
15	Numerical: For $h_t = (1 - z_t)h_{t-1} + z_t \tilde{h}_t$, given $z_t = 0.7$, $h_{t-1} = 0.4$, and $\tilde{h}_t = 0.8$, calculate h_t .	BTL-2	Understanding	CO 3
16	State the basic principle of a Bidirectional RNN (BiRNN).	BTL-1	Remembering	CO 3
17	Explain how a BiRNN uses information from both past and	BTL-2	Understanding	CO 3

	future time steps.			
18	State the meaning of a Bidirectional LSTM (BiLSTM).	BTL-1	Remembering	CO 3
19	Explain why BiLSTM is useful when both preceding and succeeding context are important.	BTL-2	Understanding	CO 3
20	Differentiate between LSTM and GRU based on their gating mechanisms and memory structure.	BTL-2	Understanding	CO 3
21	Differentiate between BiRNN and BiLSTM based on their ability to handle long-term dependencies.	BTL-2	Understanding	CO 3
22	Numerical: During BPTT, a gradient is multiplied by 0.5 at every time step. If the initial gradient is 1, calculate its value after 4 time steps.	BTL-2	Understanding	CO 3
23	Case Study: A sentiment-analysis system must understand a word using both the preceding and succeeding words. Suggest a suitable recurrent architecture and give one reason.	BTL-2	Understanding	CO 3
24	Case Study: A medical time-series application needs to learn long-term dependencies in patient observations. Suggest either a basic RNN or LSTM and justify your choice.	BTL-2	Understanding	CO 3
25	Case Study: A sequence-processing application has limited computational resources but requires gated memory. Suggest whether an LSTM or GRU would be more suitable and give one reason.	BTL-2	Understanding	CO 3
PART –B				
1	Illustrate the architecture and working principle of a Recurrent Neural Network (RNN) with a neat diagram. Demonstrate how the hidden state enables sequence processing.	BTL-3	Applying	CO 3
2	Analyze RNNs and Feedforward Neural Networks with respect to memory, sequence processing, information flow, advantages, limitations, and applications.	BTL-4	Analyzing	CO 3
3	Describe the Backpropagation Through Time (BPTT) algorithm with a neat flowchart. Derive the gradient calculation and weight-update process for a recurrent network.	BTL-3	Applying	CO 3
4	Analyze the vanishing-gradient and exploding-gradient problems encountered during BPTT. Examine suitable techniques such as gradient clipping and gated architectures to overcome them.	BTL-4	Analyzing	CO 3
5	Illustrate the architecture and working of an LSTM network with a neat diagram. Demonstrate the functions of the Forget Gate, Input Gate, Output Gate, and Cell State.	BTL-3	Applying	CO 3
6	Numerical Problem: Given the Forget Gate, Input Gate, candidate cell-state, Output Gate, previous cell state, and previous hidden state values of an LSTM cell, calculate the updated cell state and hidden state after one time step. Show all intermediate calculations.	BTL-3	Applying	CO 3
7	Analyze how LSTM overcomes the long-term dependency problem of conventional RNNs. Compare the information	BTL-4	Analyzing	CO 3

	flow and memory mechanism of LSTM with a standard RNN.			
8	Illustrate the architecture and working of a Gated Recurrent Unit (GRU) with a neat diagram. Demonstrate the roles of the Update Gate and Reset Gate.	BTL-3	Applying	CO 3
9	Analyze GRU and LSTM with respect to gating mechanism, memory representation, number of parameters, computational complexity, and ability to handle long-term dependencies.	BTL-4	Analyzing	CO 3
10	Numerical Problem: For a simplified GRU equation $h_t = (1 - z_t)h_{t-1} + z_t\tilde{h}_t$, given z_t , h_{t-1} , and $\tilde{h}_t$, calculate the updated hidden state h_t .	BTL-3	Applying	CO 3
11	Evaluate the suitability of RNN, LSTM, and GRU for speech recognition, language translation, and time-series forecasting. Justify the most appropriate architecture for each application.	BTL-5	Evaluating	CO 3
12	Illustrate the architecture and working principle of a Bidirectional RNN (BiRNN) with a neat diagram. Demonstrate how forward and backward hidden states are combined.	BTL-3	Applying	CO 3
13	Analyze the architecture and advantages of Bidirectional LSTM (BiLSTM) over conventional LSTM. Examine how BiLSTM utilizes information from both past and future contexts.	BTL-4	Analyzing	CO 3
14	Evaluate RNN, LSTM, GRU, BiRNN, and BiLSTM with respect to memory capability, sequence context, computational cost, long-term dependency handling, advantages, limitations, and applications.	BTL-5	Evaluating	CO 3
15	Case Study: A speech-recognition system fails to capture long-range dependencies in spoken sentences. Evaluate LSTM and GRU architectures and recommend the more suitable model with justification.	BTL-5	Evaluating	CO 3
16	(a) Describe the BPTT algorithm with a neat flowchart and demonstrate its application to a simple recurrent network. (8 Marks) (b) Analyze the vanishing and exploding gradient problems during BPTT and examine appropriate solutions. (8 Marks)	BTL-3 / BTL-4	Applying / Analyzing	CO 3
17	(a) Evaluate the suitability of RNN, LSTM, and GRU for a sentiment-analysis or sequence-prediction application. (8 Marks) (b) Design a suitable recurrent architecture for the selected application and justify the choice of recurrent unit, training strategy, and model configuration. (8 Marks)	BTL-5 / BTL-6	Evaluating / Creating	CO 3
18	(a) Design a recurrent neural network architecture for machine translation by selecting a suitable RNN/LSTM/GRU/BiLSTM architecture. (8 Marks) (b) Formulate the training methodology using BPTT and justify the selected architecture, sequence-processing strategy, and optimization approach. (8 Marks)	BTL-6	Creating	CO3

UNIT IV DEEP NEURAL NETWORKS				
Introduction, Difficulty of training deep neural networks, Greedy layer wise training. Generative models: Restrictive Boltzmann Machines (RBMs), Introduction to MCMC and Gibbs Sampling, gradient computations in RBMs, Deep Boltzmann Machines. Convolutional Neural Networks: LeNet, AlexNet, ZF-Net, VGGNet, GoogLeNet, ResNet, Visualizing Convolutional Neural Networks, Guided Back propagation, Deep Dream, Deep Art, Fooling Convolutional Neural Networks. Auto Encoders, Deep Reinforcement Learning, Deep Learning Tools: Caffe, Theano and Torch.				
Q.No.	Question	Bloom's Taxonomy Level	Competence	Course Outcome
PART – A				
1.	State Deep Neural Network.	BTL-1	Remembering	CO 4
2.	Identify two reasons for preferring deep neural networks over shallow networks.	BTL-1	Remembering	CO 4
3.	List the major difficulties encountered while training deep neural networks.	BTL-1	Remembering	CO 4
4.	Distinguish between vanishing-gradient and exploding-gradient problems in deep neural networks.	BTL-2	Understanding	CO 4
5.	Summarize the Greedy Layer-wise Training approach and state one advantage.	BTL-2	Understanding	CO 4
6.	State the purpose of a Generative Model in deep learning.	BTL-1	Remembering	CO 4
7.	Identify the structural restriction that distinguishes a Restricted Boltzmann Machine (RBM) from a general Boltzmann Machine.	BTL-1	Remembering	CO 4
8.	State the energy function of an RBM and identify its main components.	BTL-1	Remembering	CO 4
9.	Describe how Gibbs Sampling is used during RBM training.	BTL-2	Understanding	CO 4
10.	Differentiate between an RBM and a Deep Boltzmann Machine (DBM) based on their layer structures.	BTL-2	Understanding	CO 4
11.	State the primary purpose of a Convolutional Neural Network (CNN) in image processing.	BTL-1	Remembering	CO 4
12.	Differentiate the roles of convolution and pooling operations in a CNN.	BTL-2	Understanding	CO 4
13.	Identify two key architectural features of LeNet and AlexNet.	BTL-1	Remembering	CO 4
14.	Summarize the major improvements introduced by ZF-Net and VGGNet over earlier CNN architectures.	BTL-2	Understanding	CO 4
15.	State the purpose of the Inception Module in GoogLeNet and the residual connection in ResNet.	BTL-1	Remembering	CO 4
16.	Describe the purpose of Guided Backpropagation, DeepDream, and Deep Art in CNN visualization.	BTL-2	Understanding	CO 4
17.	Interpret how adversarial examples can cause a CNN to produce incorrect predictions.	BTL-2	Understanding	CO 4
18.	State the basic structure of an Autoencoder and mention two applications.	BTL-1	Remembering	CO 4
19.	Identify the roles of the agent, state, action, and reward in Deep Reinforcement Learning.	BTL-1	Remembering	CO 4
20.	Name three deep learning frameworks/tools and state one	BTL-1	Remembering	CO 4

	characteristic of each: Caffe, Theano, and Torch.			
21.	Numerical Problem: An input image of size 32×32 is convolved with a 5×5 filter using stride 1 and no padding. Calculate the spatial dimensions of the resulting feature map.	BTL-2	Understanding	CO 4
22.	Numerical Problem: A CNN layer receives an input of size $28 \times 28 \times 3$ and uses 16 filters of size 5×5 with one bias per filter. Calculate the total number of trainable parameters.	BTL-2	Understanding	CO 4
23.	Interpret why residual connections enable ResNet to train deeper networks than conventional CNNs.	BTL-2	Understanding	CO 4
24.	Differentiate between DeepDream and Deep Art based on their objectives and visual outcomes.	BTL-2	Understanding	CO 4
25.	Describe the role of Contrastive Divergence in training a Restricted Boltzmann Machine (RBM).	BTL-2	Understanding	CO 4
PART –B				
1	Analyze the major difficulties involved in training Deep Neural Networks, including vanishing/exploding gradients, overfitting, and computational cost. Evaluate suitable techniques for addressing these difficulties.	BTL-5	Evaluating	CO 4
2	Illustrate the Greedy Layer-wise Training procedure. Demonstrate how layer-by-layer pre-training can be used to initialize a Deep Belief Network.	BTL-3	Applying	CO 4
3	(a) Illustrate the architecture and working of a Restricted Boltzmann Machine (RBM) with a neat diagram. (8 Marks) (b) Derive the energy function and probability distributions of an RBM and interpret their significance. (8 Marks)	BTL-3 / BTL-4	Applying / Analyzing	CO 4
4	Describe the Markov Chain Monte Carlo (MCMC) framework and demonstrate Gibbs Sampling for an RBM using a suitable example.	BTL-3	Applying	CO 4
5	Derive the gradient expressions used for RBM learning with Contrastive Divergence. Trace the positive and negative phases involved in the parameter update.	BTL-4	Analyzing	CO 4
6	Analyze the architecture of a Deep Boltzmann Machine (DBM). Compare DBM, RBM, and Deep Belief Network (DBN) with respect to connectivity, learning strategy, and applications.	BTL-4	Analyzing	CO 4
7	Illustrate the architecture of a Convolutional Neural Network with a neat diagram. Demonstrate how convolution, activation, pooling, and fully connected layers transform an input image.	BTL-3	Applying	CO 4
8	(a) Illustrate the LeNet architecture and trace the flow of an input image through its major layers. (8 Marks) (b) Analyze the AlexNet architecture and evaluate its major innovations over LeNet. (8 Marks)	BTL-3 / BTL-5	Applying / Evaluating	CO 4
9	Compare ZF-Net and VGGNet with respect to network structure, filter design, depth, feature extraction, and	BTL-5	Evaluating	CO 4

	computational requirements. Evaluate their major improvements over earlier CNNs.			
10	Illustrate the GoogLeNet architecture with emphasis on the Inception Module. Analyze how parallel convolutional operations and dimensionality reduction improve computational efficiency.	BTL-4	Analyzing	CO 4
11	Analyze the ResNet architecture and demonstrate how residual/skip connections address the degradation problem and facilitate the training of very deep networks.	BTL-4	Analyzing	CO 4
12	Evaluate LeNet, AlexNet, VGGNet, GoogLeNet, and ResNet with respect to architecture, depth, parameter count, computational complexity, performance, and applications.	BTL-5	Evaluating	CO 4
13	Illustrate major CNN visualization techniques. Demonstrate the working of Guided Backpropagation for interpreting features responsible for a CNN prediction.	BTL-3	Applying	CO 4
14	Analyze DeepDream and DeepArt (Neural Style Transfer). Examine how intermediate CNN feature representations are exploited to generate artistic or dream-like images.	BTL-4	Analyzing	CO 4
15	Case Study – Medical Imaging: A hospital plans to develop an AI system for pneumonia detection from chest X-rays.(a) Evaluate LeNet, AlexNet, VGGNet, and ResNet and select the most suitable architecture with justification. (8 Marks) (b) Design a Guided Backpropagation-based visualization approach to support interpretation of the model's prediction. (8 Marks)	BTL-5 / BTL-6	Evaluating / Creating	CO 4
16	Case Study – Recommendation System: An e-commerce company wants to model user-item preferences from customer browsing data.(a) Design an RBM-based recommendation approach and specify its visible and hidden units. (8 Marks) (b) Formulate the training procedure using Gibbs Sampling and Contrastive Divergence. (8 Marks)	BTL-6	Creating	CO 4
17	Case Study – Fraud Detection: A startup has very few labeled fraudulent transactions.(a) Design an Autoencoder-based anomaly-detection system. (8 Marks) (b) Formulate a decision strategy using reconstruction error to distinguish normal and fraudulent transactions. (8 Marks)	BTL-6	Creating	CO 4
18	Case Study – Game Playing: A gaming company wants an AI agent to learn a new video game without labeled training data.(a) Design a Deep Reinforcement Learning framework by specifying the state, action, reward, policy, and value components. (8 Marks) (b) Evaluate Caffe, Theano, and Torch as development tools and recommend an appropriate framework with justification. (8 Marks)	BTL-5 / BTL-6	Evaluating / Creating	CO 4
UNIT V PARAMETER TUNING				

Newer optimization methods for neural networks (Adagrad, adadelta, rmsprop, adam, NAG), second order methods for training, Saddle point problem in neural networks, Regularization methods (dropout, drop connect, batch normalization).				
Q.No.	Question	Bloom's Taxonomy Level	Competence	Course Outcome
PART – A				
1	State the meaning of parameter tuning in neural networks.	BTL-1	Remembering	CO 5
2	List two difficulties faced by standard Gradient Descent while optimizing deep neural networks.	BTL-1	Remembering	CO 5
3	State the objective of an optimization algorithm in neural network training.	BTL-1	Remembering	CO 5
4	Recall the basic principle of the Adagrad optimization algorithm.	BTL-1	Remembering	CO 5
5	State one advantage and one limitation of Adagrad.	BTL-1	Remembering	CO 5
6	Differentiate between Adagrad and Adadelta based on their learning-rate adaptation.	BTL-2	Understanding	CO 5
7	Describe how RMSProp addresses the continuously decreasing learning-rate problem of Adagrad.	BTL-2	Understanding	CO 5
8	State the purpose of the moving average of squared gradients in RMSProp.	BTL-1	Remembering	CO 5
9	Identify the two moment estimates maintained by the Adam optimizer.	BTL-1	Remembering	CO 5
10	Describe the roles of the first moment and second moment in Adam.	BTL-2	Understanding	CO 5
11	Differentiate Adam and SGD with respect to learning-rate adaptation.	BTL-2	Understanding	CO 5
12	State the main idea of Nesterov Accelerated Gradient (NAG).	BTL-1	Remembering	CO 5
13	Differentiate NAG from the conventional Momentum method.	BTL-2	Understanding	CO 5
14	State the meaning of a second-order optimization method and name one example.	BTL-1	Remembering	CO 5
15	Identify the role of the Hessian matrix in second-order optimization.	BTL-1	Remembering	CO 5
16	State one reason why second-order optimization methods are computationally expensive for deep neural networks.	BTL-1	Remembering	CO 5
17	Describe the characteristics of a saddle point in a neural-network optimization landscape.	BTL-2	Understanding	CO 5
18	Differentiate a saddle point from a local minimum.	BTL-2	Understanding	CO 5
19	State the purpose of regularization in neural networks.	BTL-1	Remembering	CO 5
20	Describe how Dropout helps reduce overfitting in neural networks.	BTL-2	Understanding	CO 5
21	Differentiate Dropout and DropConnect based on what is randomly deactivated during training.	BTL-2	Understanding	CO 5
22	State the basic purpose of Batch Normalization and mention one advantage.	BTL-1	Remembering	CO 5
23	Describe how Batch Normalization improves the training of deep neural networks.	BTL-2	Understanding	CO 5

24	Numerical Problem: In Adagrad, if $g_t = 0.5$, accumulated squared gradient $G_t = 4$, and $\eta = 0.1$, calculate the magnitude of the parameter update using $\Delta\theta = -\frac{\eta g_t}{\sqrt{G_t}}$.	BTL-2	Understanding	CO 5
25	Numerical Problem: A neuron has an output of 6 before applying Dropout. If the dropout probability is $p = 0.2$, calculate the scaled output using inverted Dropout when the neuron is retained.	BTL-2	Understanding	CO 5
PART –B				
1	Analyze the limitations of Gradient Descent and SGD in training deep neural networks. Examine why adaptive optimization methods are required for efficient convergence.	BTL-4	Analyzing	CO 5
2	Derive the Adagrad parameter-update rule. Illustrate how accumulated squared gradients influence the adaptive learning rate, and examine its advantages and limitations.	BTL-3	Applying	CO 5
3	Analyze the Adadelta optimization method using its mathematical formulation. Demonstrate how it addresses the diminishing learning-rate problem of Adagrad.	BTL-4	Analyzing	CO 5
4	Derive the RMSProp update equations. Analyze how the exponential moving average of squared gradients improves learning stability in non-stationary optimization problems.	BTL-4	Analyzing	CO 5
5	Derive the complete Adam optimization procedure, including first-moment estimation, second-moment estimation, bias correction, and parameter update.	BTL-3	Applying	CO 5
6	Compare Adagrad, Adadelta, RMSProp, Adam, and NAG with respect to learning-rate adaptation, convergence, computational requirements, advantages, limitations, and suitable applications.	BTL-5	Evaluating	CO 5
7	Illustrate the working principle of Nesterov Accelerated Gradient (NAG). Analyze how its look-ahead mechanism differs from conventional Momentum and affects convergence.	BTL-4	Analyzing	CO 5
8	Derive the Newton optimization update using the Hessian matrix. Contrast second-order optimization with first-order methods in terms of convergence, computational cost, and memory requirements.	BTL-3	Applying	CO 5
9	Evaluate the practicality of second-order optimization methods for large-scale deep neural networks. Justify the use of Quasi-Newton or other approximate approaches where appropriate.	BTL-5	Evaluating	CO 5
10	Analyze the saddle-point problem in neural-network optimization. Examine why saddle points can hinder optimization in high-dimensional non-convex loss landscapes.	BTL-4	Analyzing	CO 5
11	Evaluate the effectiveness of Momentum, NAG, adaptive learning-rate methods, and noise injection for escaping saddle points during neural-network training.	BTL-5	Evaluating	CO 5

12	Illustrate the concept of regularization in deep neural networks. Analyze how regularization improves generalization and reduces overfitting.	BTL-3 / BTL-4	Applying / Analyzing	CO 5
13	Demonstrate the working of Dropout during training and inference. Analyze how randomly deactivating neurons helps reduce overfitting.	BTL-3	Applying	CO 5
14	Compare Dropout and DropConnect with respect to their working mechanisms, affected network components, computational requirements, and regularization effects.	BTL-4	Analyzing	CO 5
15	Derive the mathematical formulation of Batch Normalization. Illustrate its placement in a neural network and analyze its effect on training stability and convergence.	BTL-4	Analyzing	CO 5
16	(a) Analyze the training behavior of a deep CNN affected by slow convergence, unstable updates, and overfitting. Identify the likely optimization and regularization issues. (8 Marks)(b) Evaluate Adam, RMSProp, and NAG and recommend an appropriate optimizer along with Dropout or Batch Normalization for the given situation. Justify your selection. (8 Marks)	BTL-4 / BTL-5	Analyzing / Evaluating	CO 5
17	(a) Evaluate the suitability of Dropout, DropConnect, and Batch Normalization for improving generalization in a deep neural network. (8 Marks)(b) Design a suitable regularization and normalization strategy for an image-classification network and justify the placement and combination of the selected techniques. (8 Marks)	BTL-5 / BTL-6	Evaluating / Creating	CO 5
18	(a) Design a complete parameter-tuning strategy for a deep neural network running on limited computational hardware, considering optimizer, learning rate, batch size, normalization, and regularization. (8 Marks)(b) Formulate an optimization strategy for achieving fast convergence and good generalization. Justify the choice of optimizer and methods for handling saddle points. (8 Marks)	BTL-6	Creating	CO 5
