Y/K/TY

Reg. No: 22BDS0270

VIT®
Vellore Institute of Technology

Final Assessment Test – November 2025

Course: BCSE332L - Deep Learning
Class NBR(s): 1643/1644/1646/2432
Time: Three Hours
Slot: D1+TD1
Max. Marks: 100

- KEEPING MOBILE PHONE/ANY ELECTRONIC GADGETS, EVEN IN 'OFF' POSITION IS TREATED AS EXAM MALPRACTICE
- DON'T WRITE ANYTHING ON THE QUESTION PAPER

COs	CO Statements
CO1	Understand the methods and terminologies involved in deep neural network, differentiate the learning methods used in Deep-nets.
CO2	Identify and apply suitable deep learning approaches for given application.
CO3	Design and develop custom Deep-nets for human intuitive applications.
CO4	Design of test procedures to assess the efficiency of the developed model.
CO5	To understand the need for Reinforcement learning in real –time problems.

BL – Blooms Taxonomy Level (1 – Remember, 2 – Understand, 3 – Apply, 4 – Analyse, 5 – Evaluate, 6 – Create)

Answer ALL Questions
(10 X 10 = 100 Marks)

1. **Predicting Diabetes Using ANN** **CO3 BL6**

A healthcare analytics team is developing a predictive model to assist clinicians in identifying diabetic patients based on key clinical indicators as given below with output.

PatientID	HbA1c	BP	FastSugar	PostSugar	Output
P001	7.8	140	160	220	1

Design an ANN with one hidden layer with four neurons using ReLU activation and a sigmoid output neuron to classify diabetic status based on HbA1c, BP, fasting and postprandial sugar levels; then perform a forward pass using above patient data and initial weights to compute the output, and find the loss using mean square error function.

Note: Assume your initial weights between 0-1 (randomly).

2. For each of the use cases provided below, identify the most suitable activation function to be applied within an Artificial Neural Network (ANN). Justify your selection by explaining the underlying concepts, the functional behaviour of the activation, and its relevance to the specific task along with mathematical expression and graph. **CO2 BL4**

Use Case 1: Classifying weather conditions (e.g., sunny, rainy, and cloudy) based on sensor inputs like temperature, humidity, wind speed, and pressure. **[5]**

Use Case 2: Predicting whether a weather alert should be triggered (e.g., storm warning: yes/no) based on real-time sensor data **[5]**

3. Extract features from the given 5×5 binary image using a 3×3 convolutional kernel. Choose appropriate stride and padding values so that the final output feature map after convolution and pooling operations is of size 3×3. Show all calculations step-by-step and include the formulas used. **CO3 BL4**

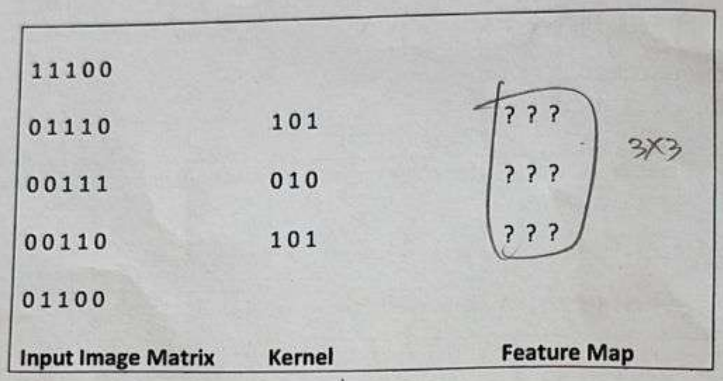

4. Compare and contrast the architectures of ResNet, AlexNet, and InceptionNet with respect to their key design features, optimization strategies, and real-world applications in deep learning. **CO3 BL3**

5. Identify and discuss a suitable optimization algorithm for training deep neural networks on MRI brain scan datasets characterized by high variability in contrast and texture. Justify your choice based on how the algorithm handles gradient stability and learning rate adaptation. **CO1 BL4**

6. In the context of handwritten digit classification, identify the type of classification problem being addressed and explain how the output activation function facilitates the prediction of digit classes. **CO1 BL3**

7.a) Given the sentence 'The movie War 2 is excellent', construct and visualize the input structure suitable for a Bidirectional RNN. Explain why BiRNN is preferred when the entire sequence is available along with a diagram. **CO2 BL3**

OR

7.b) Using BERT architecture, explain how the sentence 'The movie War 2 is excellent' is processed to generate context-aware embedding's. Describe the role of tokenization, positional encoding, and attention in forming the final context vector. **CO2 BL3**

8. A Gated Recurrent Unit (GRU) model is configured with three time steps to process a sequential input. The input sequence is given as **x = [0.5, 0.1, 0.9]**, and the initial hidden state is h0=0. For simplicity, all weight parameters are set to 1.0, and all biases are initialized to 0. The GRU cell includes update and reset gates, along with a candidate hidden state that contributes to the final hidden state at each time step. **CO4**

Using this configuration, construct the step-by-step numerical computation for each GRU cell across the three time steps. Also visualize the architecture of single GRU unit.

9.a) How do the architectural and operational differences between Single Shot Detection (SSD) and Faster R-CNN impact their suitability for latency-sensitive versus accuracy-critical applications? Discuss the implications of these trade-offs in real-world deployment scenarios. **CO3**

OR

9.b) Critically examine the architecture and learning dynamics of Generative Adversarial Networks (GANs). How does the adversarial training mechanism contribute to the generation of realistic data, and what challenges arise in stabilizing this process? **CO3**

10. In the context of intelligent decision-making systems, analyse the role of Deep Reinforcement Learning in solving complex real-time problems. Elaborate on how Deep Q-Learning integrates neural networks with Q-value estimation to enable adaptive learning in dynamic environments. **CO5**

⇔⇔⇔ **Y/K/TY** ⇔⇔⇔