Ref. No. : EX/PG/CSE/PE/T/121A/2025

M.E. COMPUTER SCIENCE & ENGINEERING 1st YEAR 2nd SEMESTER EXAMINATION 2025

DEEP LEARNING AND GENERATIVE AI

Time: 3 Hours | Full Marks: 100

This Question Paper contains two Parts: I and II. Answer each Part in a separate Answer Booklet.

Answer All sub-parts of a single Question together

PART- I [50 Marks]

Answer any one from Question Nos. 1 and 2.

Answer any two from Question Nos. 3, 4, and 5.

1. (a) What is deep learning? Compare shallow learning and deep learning. [2+4]

(b) Consider training a deep neural network with 10 layers, using the sigmoid activation function throughout. Explain in detail what happens during backpropagation in terms of gradients and how this contributes to the vanishing gradient problem. [4]

2 (a) Provide two benefits of using convolutional layers instead of fully connected ones for image analysis tasks. [2]

(b) You are benchmarking runtimes for layers commonly encountered in CNNs. Which of the following would you expect to be the fastest? Provide reason. [3]

(i) Conv layer (convolution operation + bias addition)
(ii) Max pooling
(iii) Average pooling
(iv) Batch Normalization

(c) Consider the convolutional neural network defined by the layers in the left column of the Table given in the next page. Fill in the shape of the output volume and the number of parameters at each layer. You can write the activation shapes in the format (H, W, C), where H, W, C are the height, width and channel dimensions, respectively. Unless specified, assume padding 1, stride 1, where appropriate. [5]

Notation:

- CONVx-N denotes a convolutional layer with N filters with height and width equal to x.
- POOL-n denotes a n×n max-pooling layer with stride of n and 0 padding.
- FLATTEN flattens its inputs.
- FC-N denotes a fully-connected layer with N neurons.

Layer	Activation Volume Dimension	Number of Parameters
Input	32x32x3	0
CONV3-8		
Leaky ReLU		
POOL-2		
BATCHNORM		
CONV3-16		
Leaky ReLU		
POOL-2		
FLATTEN		
FC-10		

3. (a) Your deep neural network is overfitting badly on a small dataset of medical images. You plan to apply Dropout, Batch Normalization, and L1/L2 regularization. Discuss how these three regularizations individually address overfitting. Be precise in describing the *mechanism* by which each method helps reduce overfitting. [7]

   (b) In a dataset where labels are noisy (some labels are incorrect), justify which form of regularization (L1 or L2) would be preferable and why. [3]

   (c) Explain whether switching from Stochastic Gradient Descent to Adam Optimization would meaningfully help reduce overfitting behavior. [3]

   (d) Discuss the advantages of using Mini-Batches over Stochastic Gradient Descent. [3]

   (e) Compare the effectiveness of Transfer Learning for ResNet vs. AlexNet on a medical image classification task where labeled data is scarce. [4]

4. (a) Discuss the role of Stride, Pooling and Padding in affecting the receptive field and spatial resolution. [4]

   (b) Compare the performance of Sigmoid, ReLU and Tanh activations in terms of gradient flow and convergence speed in deep networks. [6]

   (c) What is Softmax activation function? Provide a scenario when Softmax function will be used in the output layer. [4]

   (d) Compare the design philosophy of LeNet, AlexNet, and ResNet in terms of depth, receptive field expansion, and ability to capture hierarchical features. [6]

5. (a) Compare feedforward ANN and Recurrent Neural Network (RNN) in the context of sequence learning, with a focus on temporal dependencies. [4]

   (b) In the LSTM, what role does the "forget gate" play in controlling long-term memory? [2]

   (c) Explain the vanishing gradient problem for a vanilla RNN. Show, through the equations of LSTM gates, how LSTM reduces this issue. [4+6]

   (d) Discuss the advantages and limitations of Bi-directional LSTM over unidirectional LSTM.[4]

PART - II [50 MARKS]

(Use Separate Answer Booklet)

Answer any **one question from Group A** and any two questions from Group B

Group A

*Answer **any one question** from this group*

6.a) What is the main difference between generative and discriminative models? Give two examples of generative models and two examples of discriminative models.

b) Is Naive Bayes a generative or discriminative model? Justify with reasons. (3 + 2) + 5 = **10 marks**

7. a) What is the primary goal of an autoencoder?
Describe the structure of a basic autoencoder with a suitable diagram.
What is the role of the bottleneck (latent space) in an autoencoder?
Explain how image is generated using autoencoder. 2 + 3+2 + 3 =**10 marks**

Group B

*Answer **any two questions** from this group*

8. a) What is a variational autoencoder (VAE), and how does it differ from a standard autoencoder? Why is randomness introduced in the encoder of a VAE? Explain the purpose of the Kullback-Leibler (KL) divergence in VAE training.

b) Describe a GAN model with a diagram.
Write the objective function and explain how generator and discriminator of a GAN model are trained to optimize the objective function.

(2 + 3 +2 +5) +(4 + 4) = **20 marks**

9. a) Describe with a diagram a RNN-based language model. How is the loss function designed for this model. Explain how is it used for generating texts.

b) Explain with an small example self-attention mechanism used in a transformer. How does it differ from the attention mechanism used in a standard LSTM-based encoder-decoder model for text generation.

c) What is the purpose of using positional embeddings in a transformer-explain. Write the formulae used for obtaining this type of embeddings and briefly explain the intuition behind it.

(5 + 2+ 3) + (3 + 2) + (2 + 3) = **20 marks**

10. a) What is prompt Engineering? Explain with examples the following prompt Engineering techniques: i) Zero-shot ii) Few-shot iii) Chain-of-thought.

b) Describe the ethical implications of large language models in AI.

c) What is conditional GAN? Give the architecture of a conditional GAN. Give two suitable applications of conditional GAN.

(1+2 x 3) + 4 + (1+5+3) = **20 marks**
