PVP 23

Code: 23AM3601, 23DS3601

III B.Tech - II Semester - Regular Examinations - APRIL 2026

DEEP LEARNING

(Common for AIML, DS)

Duration: 3 hours Max. Marks: 70

Note: 1. This question paper contains two Parts A and B.
2. Part-A contains 10 short answer questions. Each Question carries 2 Marks.
3. Part-B contains 5 essay questions with an internal choice from each unit. Each Question carries 10 marks.
4. All parts of Question paper must be answered in one place.

BL – Blooms Level CO – Course Outcome

PART – A

		BL	CO
1.a)	Differentiate Overfitting and Underfitting in the context of deep neural networks.	L2	CO1
1.b)	Explain the purpose of a Loss Function in neural network training.	L2	CO1
1.c)	Explain Pooling Layer in a CNN.	L2	CO1
1.d)	Explain Stride and Padding in CNN.	L2	CO1
1.e)	Explain about pre-trained models. Give two examples.	L2	CO1
1.f)	Explain skip connection property in ResNet.	L2	CO1
1.g)	State two advantages of Bi-RNN over standard RNN.	L2	CO1
1.h)	Name the gates present in an LSTM cell.	L2	CO1
1.i)	List the types of GANs.	L2	CO1
1.j)	Define Generator and Discriminator in a GAN architecture.	L2	CO1

PART – B

			BL	CO	Max. Marks
UNIT-I					
2	a)	Describe the concept of Hyper parameter tuning.	L2	CO1	5 M
	b)	Explain types of activation functions commonly used in neural networks.	L2	CO1	5 M
OR					
3	a)	Explain the differences between Machine Learning and Deep Learning.	L2	CO1	5 M
	b)	Explain L1 and L2 regularization techniques in deep learning.	L2	CO1	5 M
UNIT-II					
4	a)	Explain CNN architecture with diagram.	L2	CO1	5 M
	b)	Compare Artificial Neural Network and Convolutional Neural Network.	L4	CO4	5 M
OR					
5	a)	Given input feature map dimensions 39x39 in CNN and 3 convolutional layers Layer-1: Filter Size – 3x3, Stride – 1, Padding – 0 Layer-2: Filter Size – 5x5, Stride – 2, Padding – 0 Layer-3: Filter Size – 5x5, Stride – 2, Padding – 0 Calculate output dimensions for each layer.	L3	CO2	5 M

	b)	i. Explain the role of the Fully Connected (FC) Layer in CNN. ii. Explain the role of pooling layer in CNN.	L2	CO1	5 M
UNIT-III					
6	a)	Analyze the concept of Inception modules in GoogLeNet.	L4	CO4	5 M
	b)	Explain the concept of fine-tuning and transfer learning.	L2	CO1	5 M
OR					
7	a)	Analyze the advantages of pre-trained models in different scenarios.	L4	CO4	5 M
	b)	Compare and contrast the following architectures AlexNet and VGGNet.	L4	CO4	5 M
UNIT-IV					
8	a)	Analyze the roles of input gate, forget gate and output gate in Long Short Term Memory (LSTM).	L4	CO4	5 M
	b)	Analyze the advantages of Long Short Term Memory (LSTM) over traditional Recurrent Neural Network(RNN).	L4	CO4	5 M
OR					
9	a)	Analyze the advantages of Bidirectional LSTM.	L4	CO4	5 M
	b)	Analyze the limitations of LSTMs in modern Deep Learning.	L4	CO4	5 M

UNIT-V					
10	a)	Explain about Conditional GAN.	L2	CO1	5 M
	b)	Compare and Contrast Vanilla GAN and DCGAN.	L4	CO4	5 M
OR					
11	a)	Compare GAN and Autoencoder.	L4	CO4	5 M
	b)	Explain the challenges in GAN training.	L2	CO1	5 M

III B.Tech – II Semester – Regular Examinations – APRIL 2026
DEEP LEARNING - Code: 23AM3601, 23DS3601

(Common for AIML and DS)

SHORT SCHEME OF EVALUATION

Max. Marks: 70

PART – A

Q.No	Question	BL	CO
1.a	Differentiate Overfitting and Underfitting in the context of deep neural networks. **Any two differences between overfitting and underfitting – 2M**	L2	CO1
1.b	Explain the purpose of a Loss Function in neural network training. **Explanation about Loss Function – 2M**	L2	CO1
1.c	Explain Pooling Layer in a CNN. **Pooling Layer explanation – 2M**	L2	CO1
1.d	Explain Stride and Padding in CNN. **Stride and Padding explanation – 2M**	L2	CO1
1.e	Explain about pre-trained models. Give two examples. **Pre-Trained Models explanation – 2M**	L2	CO1
1.f	Explain skip connection property in ResNet. **Skip connection property explanation – 2M**	L2	CO1
1.g	State two advantages of Bi-RNN over standard RNN. **Stating two advantages– 2M**	L2	CO1
1.h	Name the gates present in an LSTM cell. **Mentioning the names of gates – 2M**	L2	CO1
1.i	List the types of GANs. **Mentioning the List of names any two types of GANs – 2M**	L2	CO1
1.j	Define Generator and Discriminator in a GAN architecture.	L2	CO1

	Definitions of Generator and Discriminator – 2M		

PART – B

UNIT - I

Q.No	Question	BL	CO	Marks
2.a	Describe the concept of Hyperparameter tuning. **Explanation about Hyper Parameter Tuning – 5M**	L2	CO1	5 M
2.b	Explain types of activation functions commonly used in neural networks. **Explaining Any Three Types of Activation Functions – 5M**	L2	CO1	5 M

OR

Q.No	Question	BL	CO	Marks
3.a	Explain the differences between Machine Learning and Deep Learning. **Explaining Any Three Differences between Machine Learning and Deep Learning – 5M**	L2	CO1	5 M
3.b	Explain L1 and L2 regularization techniques in Deep Learning. **Explanation of L1 and L2 regularization techniques – 5M**	L2	CO1	5 M

UNIT – II

Q.No	Question	BL	CO	Marks
4.a	Explain CNN architecture with diagram. **CNN explanation – 4M** **Diagram – 1M**	L2	CO1	5 M
4.b	Compare and Contrast Artificial Neural Network and Convolutional Neural Networks **Describing Any Three Comparisons – 5M**	L4	CO4	5 M

OR

Q.No	Question	BL	CO	Marks
5.a	Given input feature map of size 39×39 in CNN and 3 convolutional layers with specified filter sizes, strides, and padding. Layer-1 : Filter Size – 3x3, Stride – 1, Padding - 0 Layer-2 : Filter Size – 5x5, Stride – 2, Padding - 0 Layer-3 : Filter Size – 5x5, Stride – 2, Padding - 0 Calculate output dimensions for each layer. **Writing the formula for calculating output dimensions – 2M** **Layer 1 output dimensions calculation – 1M** **Layer 2 output dimensions calculation – 1M** **Layer 3 output dimensions calculation – 1M**	L3	CO2	5 M
5.b	i. Explain the role of Fully Connected (FC) layer in CNN ii. Explain the role of pooling layer in CNN. **Explanation of Fully Connected Layer – 2M** **Explanation of Pooling Layer – 3M**	L2	CO1	5 M

UNIT – III

Q.No	Question	BL	CO	Marks
6.a	Analyze the concept of Inception modules in GoogleNet. **Explanation of Inception Modules in GoogleNet – 5M**	L4	CO4	5 M
6.b	Explain the concept of fine-tuning and transfer learning. **Explanation about fine-tuning and transfer learning – 5M**	L2	CO1	5 M

OR

Q.No	Question	BL	CO	Marks
7.a	Analyze the advantages of pre-trained models in different scenarios. **Explaining Any Three advantages – 5M**	L4	CO4	5 M
7.b	Compare and contrast AlexNet and VGGNet architectures. **Describing Any Two comparisons – 5M**	L4	CO4	5 M

UNIT – IV

Q.No	Question	BL	CO	Marks
8.a	Analyze the roles of input gate, forget gate and output gate in LSTM. **Explanation of Three gates – 5M**	L4	CO4	5 M
8.b	Analyze the advantages of LSTM over traditional RNN. **Explaining Any Two Advantages – 5M**	L4	CO4	5 M

OR

Q.No	Question	BL	CO	Marks
9.a	Analyze the advantages of Bidirectional LSTM. **Explaining Any Two Advantages – 5M**	L4	CO4	5 M
9.b	Analyze the limitations of LSTMs in modern deep learning. **Describing Any Two Limitations – 5M**	L4	CO4	5 M

UNIT – V

Q.No	Question	BL	CO	Marks
10.a	Explain about Conditional GAN. **Conditional GAN Explanation – 5M**	L2	CO1	5 M
10.b	Compare and contrast Vanilla GAN and DCGAN. **Describing Any Two comparisons – 5M**	L4	CO4	5 M

OR

Q.No	Question	BL	CO	Marks
11.a	Compare GAN and Autoencoder. **Describing Any Two comparisons – 5M**	L4	CO4	5 M
11.b	Explain the challenges in GAN training. **Explaining Any Two Limitations – 5M**	L2	CO1	5 M

III B.Tech – II Semester – Regular Examinations – APRIL 2026

DEEP LEARNING

(Common for AIML and DS)

DETAILED SCHEME OF EVALUATION

PART – A

Q.No	Question	BL	CO
1.a)	Differentiate Overfitting and Underfitting in the context of deep neural networks **Overfitting:** Model learns training data too well (memorization)/over training **Underfitting:** Model is too simple to learn patterns. Performs poorly on both training and testing data.	L2	CO1
1.b)	Explain the purpose of a Loss Function in neural network training • Loss function in neural network training is used calculate loss occurred or error value after training.	L2	CO1
1.c)	Explain Pooling Layer in a CNN. • Pooling layer comes after convolution layer which is used to reduce dimensions/ spatial size of feature map	L2	CO1
1.d)	Explain Stride and Padding in CNN. **Stride:** The number of steps the kernel moves across the image **Padding:** Adding zeroes around the feature map for edge feature protection	L2	CO1
1.e)	Explain about pre-trained models. Give two examples. • Pre-Trained models are the Deep Learning models which are already trained on large datasets which are used for training new datasets. Ex: Can give any two examples from the pre-trained models	L2	CO1
1.f)	Explain skip connection property in ResNet. • Skip connection allows input to skip the connections among layers in ResNet	L2	CO1
1.g)	State two advantages of Bi-RNN over standard RNN. • Understands data from both the directions • Context understanding	L2	CO1
1.h)	Name the gates present in an LSTM cell. • Input Gate	L2	CO1

	• Forget Gate • Output Gate		
1.i)	List the types of GANs. • Conditional GAN • DCGAN(Deep Convolutional GAN) • Vanilla GAN	L2	CO1
1.j)	Define Generator and Discriminator in a GAN architecture. • **Generator (G):** Generates fake data • **Discriminator (D):** Distinguishes real vs fake	L2	CO1

PART – B

2.a) Describe the concept of Hyper Parameter Tuning.

Scheme: Explanation about Hyper Parameter Tuning – 5M

Ans: Hyperparameters and Hyperparameter Tuning:

- Hyperparameters are settings that control how a machine learning model learns (like learning rate, number of neurons, etc.).
- They are not learned from data but set before training the model.
- They affect the model's accuracy and performance.
- Hyperparameter tuning is the process of finding the best values for these settings.
- Types of Hyper parameters include GridSearch CV, Random Search CV and Bayesian Optimization

2.b) Explain types of activation functions commonly used in neural networks.

Scheme: Explaining Any Three Types of Activation Functions – 5M

Ans: The activation function is used to introduce non linearity into the model and accomplish more complex tasks.

1. Sigmoid activation function: The sigmoid activation function takes input and gives 0 -1 as output

2.Tanh activation function: The tanh activation function takes input and gives -1 to +1 as output

3.ReLU activation function: ReLU stands for Rectified Linear Unit. The main advantage of using the ReLU function is avoiding negative values as output

4.Leaky ReLU activation function: Leaky ReLU mitigates the dying ReLU problem by allowing a small gradient for negative inputs.

5.Softmax activation function: The softmax function can be used for multiclass classification problems. It is known that sigmoid returns values between 0 and 1

OR

3.a) Explain the differences between Machine Learning and Deep Learning.

Scheme: Explaining Any Three Differences between Machine Learning and Deep Learning – 5M

Ans:

Machine Learning	Deep Learning
Machine Learning is a superset of Deep Learning	Deep Learning is a subset of Machine Learning
Manual Feature extraction	Automatic Feature extraction
Machine Learning Systems are simple	Deep learning systems are complex.
This uses simple linear models	These use complex non linear models
This requires less data	This requires large datasets for training
Applications: Regression, Classification, Clustering	Applications: Image Processing, Speech recognition, Object detection

3.b) Explain L1 and L2 regularization techniques in Deep Learning.

Scheme: Explanation of L1 and L2 regularization techniques – 5M

Ans:

L1 Regularization:

- It adds a penalty to the loss function using the absolute value of weights.
- It is useful for reducing model complexity and feature selection.
- It can make some weights exactly zero (removes less important features).
- Regularized cost function: **Loss + $\lambda|W_i|$**

L2 Regularization:

- It adds a penalty to the loss function using the square of weights.
- It is used to reduce large weights and prevent overfitting.
- It does not make weights exactly zero, only makes them small.
- Regularized cost function: **Loss + $\lambda(W_i^2)$**

UNIT – II

4.a) Explain CNN architecture with diagram.

Scheme: CNN explanation – 4M and Diagram – 1M

Ans: Convolution Neural Network

Convolutional Neural Network (CNN) is the extended version of artificial neural networks (ANN) which is used to extract the feature from the grid-like matrix dataset.

- **Input Layer:** Takes the image as input (pixel values).
- **Convolution Layer:** Extracts features such as edges and textures using filters.
- **ReLU Activation:** Removes negative values and keeps positive values.
- **Pooling Layer:** Reduces the size of the feature map to make computation faster.
- **Flatten Layer:** Converts 2D data into a 1D vector.
- **Fully Connected Layer:** Connects all neurons to make predictions.
- **Output Layer:** Produces the final result (e.g., classification).

(Diagram)

4.b) Compare and Contrast Artificial Neural Network and Convolutional Neural Networks

Scheme: Describing Any Three Comparisons – 5M

Ans:

ANN	CNN
Uses fully connected layers with neurons	Uses convolution, pooling, and fully connected layers
Works with numerical/tabular data	Works best with image/spatial data
Needs manual feature extraction	Automatically extracts features
Has many parameters	Has fewer parameters (shared weights)
Does not preserve spatial information	Preserves spatial relationships
Less efficient	More efficient
Higher computation for large inputs	More efficient computation

OR

5.a) Given input feature map of size 39×39 in CNN and 3 convolutional layers with specified filter sizes, strides, and padding.

Layer-1 : Filter Size – 3x3, Stride – 1, Padding - 0

Layer-2 : Filter Size – 5x5, Stride – 2, Padding - 0

Layer-3 : Filter Size – 5x5, Stride – 2, Padding - 0

Calculate output dimensions for each layer.

Scheme: Writing the formula for calculating output dimensions – 2M

Layer 1 output dimensions calculation – 1M

Layer 2 output dimensions calculation – 1M

Layer 3 output dimensions calculation – 1M

Ans:

5(a) Given input feature map of size 39×39 in CNN and 3 convolutional layers with specified filter sizes, strides and padding.

Layer 1: Filter size - 3×3, Stride = 1, Padding = 0

Layer 2: Filter Size - 5×5, Stride = 2, Padding = 0

Layer 3: Filter Size - 5×5, Stride = 2, Padding = 0

Calculate output dimensions for each layer.

Solution:- Writing formula for calculating output dimensions. (Can write any formula for calculation use)

$$\frac{N + 2p - f}{S} + 1 \quad \text{OR} \quad O = \frac{I - K + 2P}{S} + 1$$

OR

$$W_{out} = \frac{W_{in} + 2P - F}{S} + 1$$

$$H_{out} = \frac{H_{in} + 2P - F}{S} + 1$$

S = Stride

P = Padding

N/I/Win/Hin = input dimensions

F/K = filter size

Layer 1:

Input : 39×39

Filtersize (F/K) = 3×3 ; Stride (S) = 1 ; Padding (P) = 0

$$\text{Output} = \frac{N + 2P - f}{S} + 1$$

$$= \frac{39 + 2(0) - 3}{1} + 1$$

$$= 39 - 3 + 1$$

$$= 37$$

$W_{out} = 37$; $H_{out} = 37$

Layer 1 : Output dimensions $= 37 \times 37$

Layer 2: Input dimensions :- 37×37

Filter size (f) $= 5 \times 5$; Stride (s) $= 2$; padding (p) $= 0$

$$\text{Output} = \frac{N + 2p - f}{s} + 1$$

$$= \frac{37 + 2(0) - 5}{2} + 1$$

$$= \frac{32}{2} + 1$$

$$= 17$$

$W_{out} = 17$; $H_{out} = 17$

Layer 2: Output dimensions $= 17 \times 17$

Layer 3: Input dimensions :- 17×17

Filter size (f) $= 5 \times 5$; Stride (s) $= 2$; Padding (P) $= 0$

$$\text{Output} = \frac{N + 2p - f}{s} + 1 = \frac{17 + 2(0) - 5}{2} + 1$$

$$= \frac{17 - 5}{2} + 1 = \frac{12}{2} + 1 = 7$$

$W_{out} = 7$; $H_{out} = 7$

Layer 3: Output Dimensions $= 7 \times 7$

Final Answers:- Layer 1: 37×37

Layer 2: 17×17

Layer 3 : 7×7

(OR) Below solution also can be accepted for 5(a)

Individual Layer Calculations

Layer 1:-

Input dimensions : 39×39

Filter size = 3×3 ; stride (s) = 1 ; padding (p) = 0

$$\text{Output} = \frac{N + 2p - f}{s} + 1$$

$$= \frac{39 + 2(0) - 3}{1} + 1 = 37$$

Hout = 37 ; Wout = 37

Layer 1 : Output ⇒ 37 × 37

Layer 2:-

Input dimensions : 39×39 ; f = 5×5 ; stride (s) = 2 ; padding (p) = 0

$$\text{Output} = \frac{N + 2p - f}{s} + 1$$

$$= \frac{39 + 2(0) - 5}{2} + 1$$

$$= \frac{34}{2} + 1$$

$$= 18 \times 18$$

Layer 3:- Input dimensions : 39×39 ; f = 5×5 ; stride (s) = 2

padding (p) = 0

$$\text{Output} = \frac{N + 2p - f}{s} + 1 = \frac{39 + 2(0) - 5}{2} + 1$$

$$= \frac{34}{2} + 1 = 18 \times 18$$

5.b) i. Explain the role of Fully Connected (FC) layer in CNN

ii. Explain the role of pooling layer in CNN.

Scheme:

Explanation of Fully Connected Layer – 2M

Explanation of Pooling Layer – 3M

Ans: Fully Connected Layer :

- Each neuron is connected to all neurons in the previous layer.
- It combines all extracted features to make the final decision or prediction (like classification).

Pooling Layer :

- It reduces the size of feature maps to make computation easier.
- It keeps only the important information and removes less useful details.
- Common types are max pooling (takes highest value) and average pooling (takes average value)

UNIT – III

6.a) Analyze the concept of Inception modules in GoogleNet.

Scheme: Explanation of Inception Modules in GoogleNet – 5M

Ans:

- It uses different filter sizes (like 1×1, 3×3, 5×5) at the same time(parallel convolutions) to capture different features.
- It helps the network learn both small and large details from images.
- It reduces computation by using 1×1 convolutions before expensive operations.
- It improves accuracy while keeping the model efficient and not too heavy.

6.b) Explain the concept of fine-tuning and transfer learning.

Scheme: Explanation about fine-tuning and transfer learning – 5M

Ans: Transfer Learning and Fine Tuning:

- It uses a pre-trained model that is already trained on a large dataset.
- It saves time and needs less training data.
- It means slightly adjusting a pre-trained model for a new task.
- Some layers are retrained while others are kept fixed.
- It improves accuracy for the specific problem.

OR

7.a) Analyze the advantages of pre-trained models in different scenarios.

Scheme: Explaining Any Three advantages – 5M

Ans: Advantages of Pretrained Models over Traditional CNNs

1. **Reduced Training Time**
   - Already trained on large datasets, so only fine-tuning is needed.
2. **Works Well with Less Data**
   - Performs better even when the dataset is small.
3. **Better Generalization**
   - Learns general features, so it works well on new data.
4. **Cost-Effective**
   - Requires less GPU power and computing resources.
5. **Easy to Use**
   - Available in libraries like TensorFlow and PyTorch for direct use.
6. **High Accuracy Models**
   - Uses strong architectures like ResNet, VGG, and Inception.
7. **Supports Transfer Learning**
   - Can be easily adapted for new tasks with minimal training.
8. **Flexible for Different Tasks**
   - Can be used for classification, detection, and segmentation.

7.b) Compare and contrast AlexNet and VGGNet architectures.

Scheme: Describing Any Two comparisons – 5M

Ans:

AlexNet	VGGNet
8 layers (5 convolutional + 3 fully connected)	16 layers (VGG16) or 19 layers (VGG19)
Uses large filters (11×11, 5×5)	Uses small filters (3×3) repeatedly

Around 60 million parameters	Around 138 million parameters (VGG16)
Uses ReLU activation	Uses ReLU activation
Max pooling after some convolution layers	Max pooling after every 2 convolution layers
Higher error rate	Lower error rate
Less computationally expensive	More computationally expensive
Faster training	Slower training
Less used for transfer learning today	Widely used for transfer learning and fine-tuning
Designed for early GPU-based deep learning	Simple and uniform deep architecture

UNIT – IV

8.a) Analyze the roles of input gate, forget gate and output gate in LSTM.

Scheme: Explanation of Three gates – 5M

Ans: Roles of Gates in LSTM (Long Short-Term Memory)

1. Input Gate

- Decides what new information should be stored in memory.
- Controls how much of the new input is added.

2. Forget Gate

- Removes unnecessary or less important information from memory.
- Keeps only useful past information for future use.

3. Output Gate

- Selects what part of the memory should be given as output.
- Helps in producing the final prediction based on stored information.

8.b) Analyze the advantages of LSTM over traditional RNN.

Scheme: Explaining Any Two Advantages – 5M

Ans:

Advantages of LSTM over RNN

- LSTM can remember long-term information better than RNN.
- LSTM solves the vanishing gradient problem faced by RNN.
- LSTM uses gates to control what to remember and forget.
- LSTM gives higher accuracy for complex sequence tasks.
- LSTM handles long sequences efficiently, while RNN struggles.
- LSTM is better for applications like NLP, speech recognition, and time series prediction.

OR

9.a) Analyze the advantages of Bidirectional LSTM.

Scheme: Explaining Any Two Advantages – 5M

Ans: Advantages of Bidirectional LSTM (BiLSTM)

- It processes data in both forward and backward directions, so it understands full context better.
- It improves accuracy in sequence tasks like language processing by using past and future information.
- It captures more meaningful patterns compared to normal LSTM.
- It reduces ambiguity in sentences by considering complete context.
- It performs well in tasks like speech recognition, text classification, and machine translation.

9.b) Analyze the limitations of LSTMs in modern deep learning.

Scheme: Describing Any Two Limitations – 5M

Ans: Limitations of LSTMs in Modern Deep Learning

- LSTMs are slow to train because they process data step by step, making them less efficient for large datasets.
- They are computationally expensive due to complex gates and memory structures.
- LSTMs struggle with very long sequences compared to newer architectures like Transformers.
- They are harder to parallelize, which reduces training speed on GPUs.
- They may not perform as well as attention-based models in modern NLP tasks.

UNIT – V

10.a) Explain about Conditional GAN.

Scheme: Conditional GAN Explanation – 5M

Ans: Conditional GAN (cGAN)

Definition:

- Conditional GAN is an extension of GAN where both generator and discriminator use extra information (like labels or conditions) to control the type of data generated.

- The **generator** takes random noise and a condition (like class label) to generate specific data
- The **discriminator** receives real/fake data along with the same condition and checks if the data matches the condition.
- Both networks are trained together: the generator tries to fool the discriminator, and the discriminator tries to correctly identify real and fake data.

10.b) Compare and contrast Vanilla GAN and DCGAN.

Scheme: Describing Any Two comparisons – 5M

Ans: Vanilla GAN:

- Uses fully connected (dense) layers.
- Works better for simple data, not images.
- Often produces low-quality and less stable results.

DCGAN:

- Uses convolutional layers (CNN).
- Designed mainly for image generation tasks.
- Produces more realistic and stable outputs.

OR

11.a) Compare GAN and Autoencoder.

Scheme: Describing Any Two comparisons – 5M

Ans:

GAN (Generative Adversarial Network):

- Has two networks: generator and discriminator.
- Used to create new realistic data like images.
- Learns through competition between two networks.

Autoencoder:

- Has two parts: encoder and decoder.
- Used for compressing and reconstructing input data.
- Learns by trying to recreate the same input as output.

11.b) Explain the challenges in GAN training.

Scheme: Explaining Any Two Limitations – 5M

Ans:

1. **Training Instability**
   GAN training can be unstable and may not converge properly.
2. **Mode Collapse**
   Generator produces limited varieties of outputs instead of diverse samples.
3. **Vanishing Gradients**
   If the discriminator becomes too strong, the generator may stop learning.
4. **Evaluation Difficulty**
   Hard to quantitatively evaluate the quality of generated data.

5. **High Computational Cost**
   Requires powerful hardware (e.g., GPUs) and long training times.
6. **Lack of Control**
   Hard to control specific features in generated outputs (unless using variants like conditional GANs).
7. **Data Dependency**
   Requires a large amount of high-quality training data.
8. **Complex architecture**