

B.E. / B.Tech. - DEGREE EXAMINATIONS, NOV / DEC 2025

Seventh Semester

Computer Science and Engineering (AIML)

20AMPC701 - GENERATIVE DEEP LEARNING

Regulations - 2020

Duration: 3 Hours

Max. Marks: 100

PART - A (MCQ) (10 × 1 = 10 Marks)

Answer ALL Questions

	Marks	K-Level	CO
1. A probabilistic generative model primarily models: (a) $p(x y)$ (b) $p(y x)$ (c) $p(x)$ (d) $p(y)$	1	K1	CO1
2. In representation learning, a good latent space should be: (a) High-dimensional and sparse (b) Disentangled and semantically meaningful (c) Fixed random vectors (d) Unrelated to inputs	1	K1	CO1
3. The ELBO in Variational Autoencoders stands for: (a) Evidence Lower Bound (b) Estimated Latent Bound (c) Expected Loss Objective (d) Error Linear Basis	1	K1	CO2
4. Which modification stabilizes GAN training by enforcing Lipschitz continuity? (a) BatchNorm (b) Dropout (c) Wasserstein loss (d) Adam optimizer	1	K1	CO2
5. CycleGAN is primarily used for: (a) Text summarization (b) Image-to-image translation without paired data (c) Audio synthesis (d) Language modeling	1	K1	CO3
6. Neural Style Transfer separates an image into: (a) Color and texture (b) Content and style (c) Frequency and amplitude (d) Edges and corners	1	K1	CO3
7. MuseGAN is an architecture for: (a) Image generation (b) Music generation using GANs (c) Text generation (d) Speech recognition	1	K1	CO4
8. World Models combine VAE + MDN-RNN + Controller; MDN stands for: (a) Maximum Density Network (b) Mixture Density Network (c) Modular Deep Network (d) Mean Distance Network	1	K1	CO4
9. Positional encoding is necessary in transformers to: (a) Improve GPU utilization (b) Introduce token order information (c) Normalize activations (d) Reduce parameters	1	K1	CO5
10. Which model is autoregressive and excels at long-form text generation? (a) BERT (b) VAE (c) GPT-2 (d) ProGAN	1	K1	CO6

PART - B (12 × 2 = 24 Marks)

Answer ALL Questions

11. Define a probabilistic generative model and give one real-world example.	2	K1	CO1
12. Explain the difference between implicit and explicit generative models.	2	K2	CO1
13. Write the encoder-decoder objective of a basic autoencoder.	2	K1	CO2
14. Briefly explain the reparameterization trick used in VAEs.	2	K2	CO2
15. Show the high-level architecture of CycleGAN and Explain its losses.	2	K2	CO3
16. What is content loss vs style loss in neural style transfer?	2	K1	CO3
17. Explain how an RNN (LSTM) can be used to generate sequences of text.	2	K2	CO4
18. Define reward shaping in RL and how it might assist music generation models.	2	K1	CO4
19. Outline the role of VAE in World Model Architecture.	2	K2	CO5
20. Show how OpenAI Gym considered important for Reinforcement learning research.	2	K2	CO5

21. Explain multi-head attention in two sentences. 2 K2 CO6
22. Contrast BERT and GPT-2 in terms of their pretraining objectives. 2 K2 CO6

PART - C (6 × 11 = 66 Marks)

Answer ALL Questions

23. a) (i) Explain the framework of generative modelling: describe density estimation, sampling and latent-variable formulations. 7 K2 CO1
- (ii) Show how maximum likelihood estimation would fit the mixture parameters (outline steps) for the given 1D mixture of Gaussian model. 4 K2 CO1
- OR**
- b) (i) Summarize the main challenges in generative modelling (e.g., evaluation, mode coverage, scalability). 7 K2 CO1
- (ii) Compare the evaluation metrics for two image generative models and discuss their pros/cons. 4 K2 CO1
24. a) Illustrate the VAE loss (ELBO) starting from $\log p(x)$ and explain each term's role in training. 11 K3 CO2
- OR**
- b) Demonstrate the GAN training for minimax game. Show how the discriminator and generator losses are defined and why gradients can vanish. 11 K3 CO2
25. a) Make use of CycleGAN's: explain cycle-consistency loss and show with equations how unpaired image translation is possible using a simple experiment (dataset, metrics) to aCycleGAN that translates sketches \leftrightarrow photos. 11 K3 CO3
- OR**
- b) Identify neural style transfer and formulate the optimization objective and the role of different CNN layers for content and style using two ways to speed up style transfer for real-time applications and discuss trade-offs. 11 K3 CO3
26. a) Illustrate the World Model architecture (VAE + MDN-RNN + Controller). Explain training sequence for each module. 11 K2 CO4
- OR**
- b) Explain how MuseGAN structures music generation (tracks, bars, binary piano-roll representation) and discuss one limitation. 11 K2 CO4
27. a) Organize the detailed idea behind Reinforcement Learning. 11 K3 CO5
- OR**
- b) Construct the World Model Architecture and analyze the VAE and MDN-RNN training. 11 K3 CO5
28. a) Identify the architectural differences between BERT and GPT-2 and comment on why each suits its downstream tasks. 11 K3 CO6
- OR**
- b) Choose two advanced generative models (e.g., StyleGAN and BigGAN). For each, describe one architectural innovation and how it improves visual quality. 11 K3 Co5