Generative Artificial Intelligence

Code: CS32212CS	Course Type: Core	Semester: II
Evolution Scheme (TA-MSE-ESE): 20-30-100	Total Marks: 150	L-T-P (Credits): 3-0-0 (3)

Pre-Requisites:

- Foundations of Deep Learning (Feedforward, CNNs, Sequence Models)
- Advanced Probability Theory (Marginalization, Bayesian Inference, Markov Chains)
- Information Theory basics (Entropy, KL Divergence, Cross-Entropy)
- Multivariable Calculus and Continuous Optimization techniques.

Course Objective:

1. To establish the probabilistic foundations of generative modelling and analyse exact likelihood based autoregressive models.
2. To formulate latent variable models, including Variational Autoencoders (VAEs) and Normalizing Flows, using variational inference and change-of-variable theorems.
3. To investigate the game-theoretic mechanics of Generative Adversarial Networks (GANs) and resolve training instabilities mathematically.
4. To explore the stochastic processes underlying Energy-Based Models (EBMs) and Denoising Diffusion Probabilistic Models (DDPMs).

Course Outcomes: At the end of the course, the student will be able to

CO-1	Formulate maximum likelihood estimation frameworks for complex data distributions and analyze autoregressive generative architectures.
CO-2	Derive the Evidence Lower Bound (ELBO) for variational inference and apply topological transformations in Normalizing Flows.
CO-3	Analyze the minimax optimization landscape of GANs, derive the Wasserstein distance metric, and implement mathematically stable adversarial networks.
CO-4	Formulate forward and reverse stochastic differential equations for diffusion models and apply score-matching techniques to generate high-fidelity samples.

Course Articulation Matrix:

PO	PEO-1	PEO-2	PEO-3	PEO-4	PEO-5
CO-1	3	2	2	1	2
CO-2	3	3	1	-	3
CO-3	3	3	2	1	3
CO-4	3	3	2	2	3

1 - Slightly; 2 - Moderately; 3 – Substantially

Syllabus:

Unit-1: Probabilistic Foundations and Autoregressive Models:

Taxonomy of generative models. Mathematical foundations: Maximum Likelihood Estimation (MLE), Kullback-Leibler (KL) divergence, and Jensen-Shannon (JS) divergence. Explicit density modeling via Autoregressive architectures: Chain rule of probability and conditional independence assumptions. Architectures: Fully Visible Sigmoid Belief Networks (FVSBN), NADE, PixelRNN, and PixelCNN. Masked convolutions and blind spots.

Sequence generation fundamentals: Transformers as autoregressive sequence models and the mathematical formulation of self-attention for generative tasks.

Unit-2: Latent Variable Models and Normalizing Flows:

Latent variable formulations and the intractability of exact inference. Variational Inference: Derivation of the Evidence Lower Bound (ELBO). Variational Autoencoders (VAEs): The reparameterization trick, encoder/decoder parameterization, and properties of the Gaussian prior. Enhancements: VAE for disentanglement, Hierarchical VAEs, and Vector Quantized VAEs (VQVAE). Normalizing Flows: The change-of-variables theorem and Jacobian determinants. Flow architectures: RealNVP, Glow, and continuous normalizing flows via Neural ODEs.

Unit-3: Generative Adversarial Networks (GANs):

Implicit density models and the game-theoretic formulation of Generative Adversarial Networks. The minimax objective, optimal discriminator derivation, and the relation to JS divergence. Training pathologies: Vanishing gradients and mode collapse. Advanced adversarial metrics: Integral Probability Metrics (IPMs) and the Earth Mover's (Wasserstein) distance. Wasserstein GAN (WGAN) and Gradient Penalty (WGAN-GP) derivations. Conditioning GANs (cGANs), Image-to- Image translation frameworks, and StyleGAN architectures.

Unit-4: Score-Based & Diffusion Probabilistic Models:

Energy-Based Models (EBMs): Formulation, partition function intractability, and training via contrastive divergence. Score matching techniques: Denoising score matching and Langevin dynamics. Denoising Diffusion Probabilistic Models (DDPMs): Forward Gaussian noising processes and reverse denoising transitions. Evidence lower bound for diffusion. Continuous-time formulations: Stochastic Differential Equations (SDEs) for score-based generative modeling (VESDE and VP-SDE). Accelerated sampling (DDIM) and Latent Diffusion Models (Stable Diffusion framework).

Learning Resources:

Text Books:

1. Understanding Deep Learning (Ch 17–21) Prince S.J.D. MIT Press 2024
2. Deep Generative Modeling Tomczak J.M. Springer 2022.
3. Deep Learning (Ch 20) Goodfellow I., Bengio Y., Courville A. MIT Press 2016

Reference Books:

3. Probabilistic Machine Learning: Advanced Topics Murphy K.P. MIT Press 2023
4. Generative Deep Learning Foster D. O'Reilly Media 2023 (2nd Ed.)