

NIRMA UNIVERSITY

Institute:	Institute of Technology
Name of Programme:	BTech CSE (Artificial Intelligence & Machine Learning)
Course Code:	XXXX
Course Title:	Generative AI
Course Type:	Department Elective-II
Year of Introduction:	2025-26

L	T	Practical Component				C
		LPW	PW	W	S	
3	0	2	-	-	-	4

Course Learning Outcomes (CLOs):

At the end of the course, the students will be able to -

1. describe foundational concepts of generative models. (BL1)
2. implement basic autoencoder-based generative models. (BL3)
3. evaluate advanced transformer and diffusion models for various generative tasks. (BL5)
4. design responsible and ethical generative AI applications for real-world use cases. (BL6)

Unit	Contents	Teaching Hours (Total 45)
Unit-I	Introduction to Generative AI: Definition, history, scope; comparison with discriminative models; applications in text, vision, speech, code.	04
Unit-II	Basic Generative Models: Gaussian mixture models, Hidden Markov models, Variational Autoencoders (VAE), encoder-decoder framework, reconstruction loss, KL divergence.	08
Unit-III	Generative Adversarial Networks (GANs): Architecture and training, loss functions, mode collapse, variants – DCGAN, CycleGAN, StyleGAN.	08
Unit-IV	Transformer Models for Generation: Attention mechanism, transformer architectures, autoregressive models, GPT & BERT family, ICL, Prompt Engineering.	11
Unit-V	Diffusion Models and Emerging Architectures: Score-based models, Denoising Diffusion Probabilistic Models (DDPM), introduction to SDEs in generative modelling.	08
Unit-VI	Ethical and Societal Issues in Generative AI: Deepfakes, misinformation, copyright, fairness, privacy, AI alignment, responsible AI practices.	06

Self-Study:

The self-study contents will be declared at the commencement of the semester. Around 10% of the questions will be asked from self-study content.

Suggested Readings/ References:

1. David Foster, Generative Deep Learning: Teaching Machines to Paint, Write, Compose, and Play, 2nd Edition, O'Reilly Media.
2. Ian Goodfellow, Yoshua Bengio, and Aaron Courville, Deep Learning, MIT Press.

3. Max Welling et al., Introduction to Variational Autoencoders, Now Publishers.
4. Kevin Murphy, Probabilistic Machine Learning: Advanced Topics, MIT Press.

Laboratory Work:

Laboratory work will be based on the above syllabus with a minimum of 10 experiments to be incorporated. The students in a suitable group size will design and perform one experiment as a part of Laboratory work.

Sr. No.	List of Experiments/Exercises	Hours
1	Getting Started with Generative AI: Install Hugging Face Transformers and Diffusers, and run a basic pre-trained model, such as GPT-2, for text generation.	4
2	Variational Autoencoder: Implement VAE on MNIST dataset using PyTorch. Visualize latent space interpolations.	4
3	GAN from Scratch: Implement a simple GAN on the MNIST dataset. Observe generator/discriminator losses and generated samples.	4
4	Style Transfer with Pretrained GANs: Use StyleGAN or CycleGAN for image translation (e.g., horse to zebra).	4
5	Prompt Engineering: Explore prompt tuning with GPT-3 or open-source alternatives like Falcon or Mistral using the API.	4
6-7	Image Generation using Diffusers: Generate images using Stable Diffusion and customize sampling parameters.	4
8-9	Multimodal Generation: Use BLIP or DALL·E for text-to-image or image captioning tasks.	4
10	Ethical AI Audit: Analyze outputs from generative models for bias, misinformation, and plagiarism.	2