

Title of the Course	Application of Deep Learning and Neural Networks
----------------------------	---

Credits	T	P	E	C
	3	1	0	4
Course Type (Theory/Practical/Integrated)	Integrated			
Course Category	Core Discipline			
Pre-Requisite	<ol style="list-style-type: none"> 1. Linear Algebra: matrices, vectors 2. Calculus: differentiation, partial derivatives, gradients 3. Statistics: normal distribution, probability 4. Basic programming constructs 5. Data processing libraries: NumPy, SciPy, Matplotlib, Pandas 6. Basic use of Scikit-learn 7. Completion of the TCS iON Machine Learning for Real-World Application course 			
Learning Objectives	<ol style="list-style-type: none"> 1. Understand the fundamentals of deep learning, neural networks, and advanced CNN architectures. 2. Analyze sequence modeling techniques using RNNs, LSTMs, GRUs, and HMMs. 3. Apply transformer-based architectures for modern NLP and generative AI tasks. 4. Evaluate generative models such as VAEs, GANs, and large language models for real-world applications. 5. Incorporate responsible AI principles while designing and deploying deep learning systems. 			
Course Outcomes & Bloom's Level	CO Code	Course Outcome Statement	Bloom Level	
	CO1	Explain the core concepts of neural networks, deep learning fundamentals, and CNN architectures.	L2	
	CO2	Apply sequence modeling techniques using RNNs, LSTMs, GRUs, and HMMs for temporal data analysis.	L3	
	CO3	Analyze transformer architectures including encoder-only, decoder-only, and encoder-decoder models.	L4	
	CO4	Evaluate the performance and use-cases of generative AI models such as VAEs, GANs, and LLMs.	L5	
	CO5	Design deep learning solutions using modern frameworks and libraries like Hugging Face.	L6	

	CO6	Assess ethical considerations and implement responsible AI practices in deep learning systems.	L5
Course Elements	Course Element	Coverage Level	
	Skill Development	✓✓✓ (High)	
	Entrepreneurship	✓✓✓ (High)	
	Employability	✓✓ (Moderate)	
	Professional Ethics	✓✓✓ (High)	
	Gender	✓✓ (Moderate)	
	Human Values	✓ (Low/Indirect)	
	Environment & Sustainability	✓ (Low/Indirect)	
SDG (Goals)	SDG 4: Quality Education		
Total Hours of Pedagogy	45 hours Theory (15 hours Self-paced content + 30 hours lecture) 30 hours Practical		

Module#	Content	Pedagogy
M-1	Deep Learning Fundamentals & Advanced CNNs Introduction to Neural Networks: Perceptron, Feed-Forward Neural Networks; Deep Learning Frameworks: PyTorch vs. TensorFlow/Keras – choosing the right tool; Advanced CNN Architectures: ResNet, Inception, MobileNet; Object Detection: R-CNN family (Faster R-CNN, Mask R-CNN), YOLO, SSD.	Self-paced content, Lecture, Modular Assignment
M-2	Recurrent Neural Networks and Sequence Modelling Hidden Markov Models; RNNs Deep Dive: Fundamentals, Vanishing/Exploding Gradients; Long Short-Term Memory (LSTM) Networks: Architecture, gates; Gated Recurrent Units (GRUs): Simplified recurrent units; Sequence-to-Sequence Models: Encoder-Decoder architecture for machine translation; Attention Mechanisms: Precursor to Transformers; Applications: Advanced Time Series Forecasting (using DL), Speech Recognition.	Self-paced content, Lecture, Modular Assignment
M-3	The Transformer Architecture The Transformer Architecture: Self-attention, Multi-head attention, Positional Encoding; Encoder-only Transformers: BERT, RoBERTa, Electra – Masked Language Modelling, Next Sentence Prediction; Decoder-only Transformers: GPT series, Llama, Mistral – Autoregressive generation; Encoder-Decoder Transformers: T5, BART – Unified text-to-text framework; Hugging Face Transformers Library: Model loading, tokenizers, pipelines, fine-tuning.	Self-paced content, Lecture, Modular Assignment
M-4	Generative AI Models Autoencoders; Variational Autoencoders (VAEs): ELBO derivation, training stable VAEs, conditional VAEs; Generative Adversarial Networks (GANs): Architecture, Training Stability Issues: Mode collapse, non-convergence, Conditional GANs (cGANs), Applications: Photo-realistic image synthesis, image-to-image	Self-paced content, Lecture, Modular Assignment

	translation, data augmentation; Advanced Large Language Models (LLMs) - Introduction to LLMs – Language Models, LLM Pre-training and Fine-tuning Strategies: Full fine-tuning vs. Parameter-Efficient Fine-Tuning (PEFT) like LoRA, QLoRA. Quantization techniques, Learning from Human Feedback (RLHF): Aligning LLMs with human preferences, techniques are Reinforcement Learning from Human Feedback (RLHF), Direct preference optimization (DPO), etc., Retrieval-Augmented Generation (RAG): Architecture, vector databases, chunking, embedding models, implementing RAG, LLM Evaluation: Quantitative metrics (Perplexity, BLEU, ROUGE), qualitative evaluation, human-in-the-loop, LLM-as-a-Judge; Multimodal AI: Combining text, image, and other modalities for unified generative tasks.	
M-5	Responsible AI Responsible AI in Deep Learning - Bias Detection and Mitigation in DL: Datasets bias, algorithmic bias, post-hoc methods; Explainable AI (XAI) for DL: LIME, SHAP, Grad-CAM, attention visualization.	Self-paced content, Lecture, Modular Assignment

List of DIY Modular Assignments

1. Deep Learning Fundamentals and Advanced CNNs
2. Semantic Segmentation with U Net (Oxford IIIT Pet)
3. Building forecasters and dataset handling
4. EN to DE Translation and BLEU Comparison
5. Fine-Tuning BERT/DistilBERT for AG News Classification
6. Encoder–Decoder Summarization (T5 on SAMSum)
7. DCGAN on Fashion-MNIST
8. Grad-CAM + Responsible AI (CIFAR-10)

Modular Assignment Mapping

S. No.	DIY Assignment Title	Mapped CO(s)	Bloom Level
1	Deep Learning Fundamentals and Advanced CNNs	CO1	L2
2	Semantic Segmentation with U-Net (Oxford IIIT Pet)	CO5	L6
3	Building Forecasters and Dataset Handling	CO2	L3
4	EN to DE Translation and BLEU Comparison	CO3	L4
5	Fine-tuning BERT/DistilBERT for AG News Classification	CO5	L6
6	Encoder-Decoder Summarization (T5 on SAMSum)	CO3, CO5	L5
7	DCGAN on Fashion-MNIST	CO4	L5
8	Grad-CAM + Responsible AI (CIFAR-10)	CO6	L5

Reference Books	<ol style="list-style-type: none"> 1. Ian Goodfellow, Yoshua Bengio, Aaron Courville, <i>Deep Learning</i>, 1st Edition, MIT Press, 2016. 2. Francois Chollet, <i>Deep Learning with Python</i>, 2nd Edition, Manning Publications, 2021.
------------------------	---

	<ol style="list-style-type: none"> 3. Aurelien Geron, <i>Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow</i>, 3rd Edition, O'Reilly Media, 2022. 4. Simon Haykin, <i>Neural Networks and Learning Machines</i>, 3rd Edition, Pearson Education, 2009. 5. Charu C Aggarwal, <i>Neural Networks and Deep Learning: A Textbook</i>, 1st Edition, Springer, 2018.
--	---

Course Articulation Matrix

CO \ PO	PO1	PO2	PO3	PO4	PO5	PO6	PO7	PO8	PO9	PO10	PO11	PO12
CO1	3	2	–	–	2	–	–	–	–	–	–	1
CO2	2	3	2	2	3	–	–	–	–	–	–	2
CO3	2	3	2	3	3	–	–	–	–	–	–	2
CO4	2	3	1	3	2	–	–	–	–	–	–	2
CO5	2	2	3	3	3	–	–	–	1	1	1	2
CO6	–	2	–	2	–	2	1	3	1	–	–	2