

25MCST401	DEEP LEARNING	Category	L	T	P	R	Credit	Year of Introduction
		VAC	4	0	0	0	4	2025

Preamble: This course explores advanced neural network architectures including Convolutional Neural Networks (CNNs) for image-related applications and Recurrent Neural Networks (RNNs) for sequential and text data processing and introduces autoencoders and generative models, which are widely used for representation learning and data generation.

Prerequisite: Nil

Course Outcomes: After the completion of the course, the student will be able to:

CO	Course Outcomes	Knowledge Level
CO1	Apply deep neural networks using suitable activation functions for pattern recognition and classification.	K3 Apply
CO2	Apply CNN and RNN models, for solving image and sequence learning problems.	K3 Apply
CO3	Apply the concepts of Recurrent Neural Network, its variants and familiarize natural language processing fundamentals.	K3 Apply
CO4	Apply the concepts of auto encoder, generative models for advanced AI operations.	K3 Apply

Mapping of course outcomes with program outcomes and program specific outcomes

	PO1	PO2	PO3	PO4	PO5	PO6	PO7	PO8	PO9	PO10	PO11	PSO1	PSO2
CO1	3	3	3	3	-	-	-	-	-	-	3	3	2
CO2	3	3	3	3	3	-	-	-	-	-	3	3	2
CO3	3	3	3	3	3	-	-	-	-	-	3	3	3
CO4	3	3	3	3	-	-	-	-	-	-	3	3	3
Avg	3	3	3	3	3	-	-	-	-	-	3	3	2.5

Mapping of course outcomes with Knowledge and Attitude Profile (WK)

	WK1	WK2	WK3	WK4	WK5	WK6	WK7	WK8	WK9
CO1	*	*	*	*	*			*	
CO2	*	*	*	*	*	*		*	
CO3	*	*	*	*	*	*		*	
CO4	*	*	*	*	*			*	

Wks Mapped with the course

WK1	WK2	WK3	WK4	WK5	WK6	WK7	WK8	WK9
*	*	*	*	*	*		*	

Mark Distribution:

Continuous Internal Evaluation Pattern:

Total Marks	CIE Marks	ESE Marks	ESE Duration
100	40	60	2.5 hrs

Attendance : 5 marks
Continuous Assessment Test : 20 marks
Continuous Assessment Assignment : 15 marks

SYLLABUS

Module 1 (14 hrs)

Introduction, Deep feed forward network. Activation functions – Sign, Sigmoid, Tanh, ReLU, leaky ReLU, Hard Tanh, Softmax. Loss function. Training a Neural Network with Backpropagation. Practical issues in neural network training. Overfitting, Underfitting, Hyper parameters and Validation sets, Estimators -Bias and Variance.

Module 2 (14 hrs)

Convolutional Neural Networks (CNNs): Motivation, Architectural Overview, CNN Layers, Filters/Kernels, Parameter Sharing, Regularization Techniques. Popular CNN Architectures: AlexNet, ResNet. Applications of CNNs in image classification, object detection, face recognition, and computer vision.

Module 3 (14 hrs)

Recurrent Neural Networks (RNNs): RNN Architecture, Training using Backpropagation Through Time (BPTT), Bidirectional Recurrent Neural Networks (Bi-RNNs), Encoder–Decoder and Sequence-to-Sequence Architectures, Long Short-Term Memory (LSTM) Networks., and Gated Recurrent Units (GRU).

Module 4 (14 hrs)

Auto Encoders and Generative models. Autoencoders- Variational AutoEncoder-under complete Auto-encoder, stochastic encoder, denoising encoder; Applications of Autoencoders. Generative models - Boltzmann machines- Deep Belief Networks- Generative Adversarial Networks

Text Books:

1. Charu C. Aggarwal. (2023). *Neural Networks and Deep Learning* (2nd ed.). Springer.
2. Ian Goodfellow, Yoshua Bengio and Aaron Courville. (2016). *Deep Learning*. MIT Press.

Reference Books:

1. François Chollet and Matthew Watson, (2025). *Deep Learning with Python* (3rd ed.). Manning Publications.
2. Umberto Michelucci. (2018). *Applied Deep Learning: A Case-Based Approach to Understanding Deep Neural Networks*. Apress.
3. Antonio Gulli and Sujit Pal. (2017). *Deep Learning with Keras*. Packt Publishing.

END SEMESTER EXAMINATION -MODEL QUESTION PAPER

Course Code	25MCST401
Course Name	DEEP LEARNING
Max. Marks: 60	Duration: 2.5 Hours

	Part A <i>(Answer all questions. Each question carries 3 marks)</i>			
Qn No.	Question	Marks	CO	BL
1	Explain how vanishing gradients affect the training of deep neural networks.	3	1	L2
2	A neural network model achieves 99% training accuracy but only 70% testing accuracy. Apply your understanding of model evaluation to identify the issue affecting the model and explain how appropriate techniques can be used to improve its performance on unseen data.	3	1	L3
3	Explain how parameter sharing reduces the number of trainable parameters in a CNN.	3	2	L2
4	A deep convolutional neural network suffers from reduced accuracy as additional layers are added. Apply the concept of skip	3	2	L3

	connections in a ResNet architecture to address this issue and explain how they improve the training of the network.			
5	A language processing application needs to predict the sentiment of a sentence. Apply the architecture of a Recurrent Neural Network (RNN) to process the sequence of words in the sentence and explain how information is propagated through the network to generate the output.	3	3	L3
6	A sentiment analysis system must classify customer reviews as positive or negative. Apply a Bidirectional RNN to process the review text and explain how information from both past and future contexts is used to produce the final classification.	3	3	L3
7	A medical imaging system receives noisy MRI images that need to be enhanced before diagnosis. Apply a denoising autoencoder to reconstruct the original images and explain how the encoder and decoder work together to remove noise	3	4	L3
8	Explain how a Restricted Boltzmann Machine (RBM) can be applied for feature learning in a simple dataset.	3	4	L2
	Part B <i>(Answer any one full question from each Module. Each carry 9 marks)</i>			
	Module I			
Qn No.	Question	Marks	CO	BL
9.a	A neural network designed for image classification fails to converge during training. Apply appropriate modifications to the learning rate, activation functions, and weight initialization methods to improve convergence, and explain how each modification addresses the problem.	5	1	L3
9.b	A neural network is being trained to classify handwritten digits. Apply the backpropagation algorithm during training and explain how forward propagation, error calculation, gradient computation, and weight updates are performed to improve classification accuracy.	4	1	L3
	OR			
10.a	A neural network developed for image classification exhibits poor performance on unseen test data. Apply appropriate hyperparameter tuning techniques using a validation set to improve model generalization, and explain how the adjustments help balance bias and variance.	5	1	L3
10.b	A machine learning engineer is developing models for binary classification and deep image classification tasks. Apply suitable activation functions (Sign, Sigmoid, and ReLU) to different layers of the models and explain why each function is appropriate for the given task.	4	1	L3

	Module II			
Qn No.	Question	Marks	CO	BL
11.a	A CNN model developed for object recognition achieves high training accuracy but poor test accuracy, indicating overfitting. Apply appropriate regularization techniques such as dropout, data augmentation, and weight regularization to improve the model's generalization performance, and explain the role of each technique.	5	2	L3
11.b	A computer vision system is required to detect objects in images. Apply convolutional filters (kernels) in a CNN to extract edge and texture features from an input image, and explain how these features contribute to object recognition.	4	2	L3
	OR			
12.a	A hospital has a dataset of labeled tumor images and wants to develop an automated tumor classification system. Apply a Convolutional Neural Network (CNN) to this dataset and explain the steps involved in identifying and classifying tumors.	5	2	L3
12.b	A researcher is developing an image classification system for identifying different species of plants. Apply the architectural features of AlexNet, such as convolutional layers, ReLU activation, pooling, and dropout, and explain how they can be used to improve classification performance.	4	2	L3
	Module III			
Qn No.	Question	Marks	CO	BL
13.a	A weather forecasting system uses historical temperature data to predict future temperatures. Apply a Recurrent Neural Network (RNN) trained using Backpropagation Through Time (BPTT) and explain how the network updates its weights across time steps to improve prediction accuracy.	5	3	L3
13.b	A language translation system must process long sentences while retaining important contextual information. Apply an LSTM network to this task and explain how its memory cells and gates help improve performance in modeling long sequences.	4	3	L3
	OR			
14.a	A mobile keyboard application suggests the next word while a user is typing a sentence. Apply a GRU-based architecture to this next-word prediction task and explain how the update gate and reset gate help in generating accurate predictions.	5	3	L3
14.b	A news agency wants to automatically generate short summaries of lengthy news articles. Apply an Encoder–Decoder model to this text summarization task and describe the roles of the encoder and decoder in generating the summary.	4	3	L3

	Module IV			
Qn No.	Question	Marks	CO	BL
15.a	A designer wants to generate new handwritten digit images similar to an existing dataset. Apply a Variational Autoencoder (VAE) to this task and explain how stochastic encoding and latent space sampling are used to generate new images.	5	4	L3
15.b	A dataset contains a large number of features, making classification computationally expensive. Apply the encoder–decoder structure of an autoencoder to extract meaningful features and explain how the extracted features can be used to improve the performance of a classifier.	4	4	L3
	OR			
16.a	A company wants to generate realistic synthetic images of handwritten digits to augment its training dataset. Apply a Generative Adversarial Network (GAN) for this task and explain how the Generator and Discriminator interact during the image generation process. .	5	4	L3
16.b	A facial recognition system is required to classify images of different individuals. Apply a Deep Belief Network (DBN) to the given image dataset and explain how hierarchical features extracted at different layers can improve recognition accuracy.	4	4	L3