20MCA283	DEEP LEARNING	CATEGORY	L	T	P	CREDIT
		ELECTIVE	3	1	0	4

Preamble: This course intends to provide insight into deep learning. This topic is currently a much sought-after skill and is under active research. Students have to refer appropriate research papers and multiple books to get in-depth knowledge about the topics**.** Instructors may give suitable programming assignments to augment the material covered in the classroom.

Prerequisite: Basic concepts of linear algebra, probability and optimization.

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

CO No.	Course Outcome (CO)	Bloom's Category Level
CO 1	Explain the basic concepts of deep learning.	Level 2: Understand
CO 2	Design neural networks using TensorFlow	Level 3: Apply
CO 3	Solve real world problems with CNN.	Level 3: Apply
CO 4	Solve real world problems with RNN.	Level 3: Apply
CO 5	Describe the concepts of GAN.	Level 2: Understand

Mapping of course outcomes with program outcomes

	PO 1	PO 2	PO 3	PO 4	PO 5	PO 6	PO 7	PO 8	PO 9	PO 10	PO 11	PO 12
CO 1	2	2										
CO 2	3	3	3		3		3					
CO 3	3	3	3		3		3					
CO 4	3	3	3		3		3					
CO 5	2	3			2		2					

3/2/1: High/Medium/Low

Assessment Pattern

Bloom's Category	Continuous Assessment Tests		End Semester Examination
	1	2	
Remember	15	15	10
Understand	25	25	30
Apply	10	10	20
Analyze			
Evaluate			
Create			

Mark distribution

Total Marks	CIE	ESE	ESE Duration
100	40	60	3 hours

Continuous Internal Evaluation Pattern:

Attendance : 8 marks
Continuous Assessment Test (2 numbers) : 20 marks
Assignment/Quiz/Course project : 12 marks

End Semester Examination Pattern: There will be two parts; Part A and Part B. Part A contains 10 compulsory short answer questions, 2 from each module. Each question carries 3 marks. Part B contains 2 questions from each module of which student should answer any one. Each question can have maximum 2 sub-divisions and carry 6 marks

Course Level Assessment Questions

Course Outcome 1 (CO1):

1. Describe the model of a biological neuron.
2. Explain Perceptron learning algorithm.
3. Explain the role of batch normalization in training a neural network.

Course Outcome 2 (CO2)

1. Draw and demonstrate the VGG-16 architecture.
2. Sketch the AlexNet architecture and explain its functionalities.

Course Outcome 3(CO3):

1. Design a convolutional neural network which can classify MNIST handwritten data.
2. An input image has been converted into a matrix of size 12 X 12 along with a filter of size 3 X 3 with a Stride of 1. Determine the size of the convoluted matrix.
3. Why do we prefer Convolutional Neural networks (CNN) over Artificial Neural networks (ANN) for image data as input?

Course Outcome 4 (CO4):

1. You are given an image data set with 10 classes. Describe how you will use deep learning to build a classifier.
2. Design a system to generate deep fakes from an image.

Course Outcome 5 (CO5):

1. Describe auto encoders and how they help in dimensionality reduction.
2. Explain how GANS work.

Model Question Paper
Course Code: 20MCA283

Course Name: DEEP LEARNING

Max. Marks :60 **Duration: 3 Hrs**

Part A

Answer all questions. Each question carries 3 marks (10 * 3 = 30 Marks)

1. Describe sigmoid activation functions.
2. Write the gradient descent algorithm.
3. Explain with an example how graphs are stored and represented in TensorFlow.
4. Discuss how graph representation can accelerate computing models.
5. Describe the VGG 16 architecture.
6. What is max pooling in the context of CNN?
7. Explain ReLU.
8. Explain the problem of vanishing gradients.
9. Write a note on auto encoders.
10. Explain the idea behind cross entropy.

Part B

Answer one full question from each module, each carries 6 marks.

11. (a) Describe the model of a biological neuron. 3 marks

(b) Explain perceptron learning algorithm. 3 marks

OR

12. With a suitable example explain how backpropagation works 6marks

13. Explain the role of batch normalization in training a neural network and describe how to find out overfitting from training and validation curves 6 marks

OR

14. Explain the ideas of Rank, Shape and Type with an example in the context of a Tensor Data Structure 6 marks

15. With a suitable numerical example illustrate convolution operation. 6 marks

OR

16. Explain the architecture of AlexNet. 6 marks

17. Explain the idea of Truncated backpropagation through time. 6 marks

OR

18. Describe how LSTM works. 6 marks

19. Distinguish between generative and discriminative models 6 marks

OR

20. Explain how a GAN is trained. 6 marks

Syllabus

Module I (8 Hours)
Review of Neural Networks: Model of a biological neuron, McCulloch Pitts Neuron, Activation Functions, Perceptron, Perceptron Learning Algorithm and Convergence, Multilayer Perceptron, Back propagation, Learning XOR, Sigmoid Neurons, Gradient Descent, Feed forward Neural Networks.
Module II (10 Hours)
Training Neural Networks: Initialization, dropout, batch normalization and dropout, overfitting, underfitting, training and validation curves. Data Visualization: Feature and weight visualization, tSNE. Introduction to TensorFlow: graphs, nodes, Tensor data structures - rank, shape, type, Building neural networks with TensorFlow, Introduction to Keras.
Module III (10 Hours)
Convolutional Neural Networks: Convolution operation, Convolutional layers in neural network, pooling, fully connected layers. Case study: Architecture of Lenet, Alexnet and VGG 16
Module IV (8 Hours)
Recurrent Neural Networks: Back propagation, vanishing gradients, exploding gradients, truncated backpropagation through time, Gated Recurrent Units (GRUs), Long Short-Term Memory (LSTM) cells, solving the vanishing gradient problem with LSTMs.
Module V (9 Hours)
Autoencoders, variational autoencoders. Generative Adversarial Networks (GAN): Discriminative and generative models, GAN discriminator, GAN generator, upsampling, GAN Training, GAN challenges, loss functions, cross entropy, minimax loss, Wasserstein loss.

Programming assignments using TensorFlow maybe given at the end of each module to get hands on experience.

Textbooks.

1. Generative Deep Learning: David Foster, OReily, (2019)
2. Deep Learning, Ian Goodfellow, Yoshua Bengio and Aaron Courville, MIT press (2016)
3. Hands on Machine Learning with Scikit Learn and TensorFlow, Aurélien Géron (2019)
4. Deep Learning Illustrated, Jon Krohn, Grant Beyleveld, Aglae Bassens, Pearson, 1st Edn., (2020)
5. Online book Dive Deep into Machine Learning at https://d2l.ai/

References

Module 1

a. https://www.cse.iitm.ac.in/~miteshk/CS6910/Slides/Lecture2.pdf
b. https://www.cse.iitm.ac.in/~miteshk/CS6910/Slides/Lecture3.pdf

Module 2

a. http://neuralnetworksanddeeplearning.com
b. Hands on Machine Learning with Scikit Learn and TensorFlow, Aurélien Géron
c. Probabilistic Machine Learning: An Introduction, Kevin Murphy
d. https://www.researchgate.net/publication/228339739_Viualizing_data_using_t-SNE

Module 3

a. https://www.cse.iitm.ac.in/~miteshk/CS7015/Slides/Teaching/pdf/Lecture11.pdf
b. Convolutional neural networks for visual computing (Chapter 4), Ragav Venkatesan and Baoxin Li CRC press

Module 4

a. On the difficulty of training RNNs: https://arxiv.org/pdf/1211.5063.pdf
b. LSTM: A Search Space Odyssey: https://arxiv.org/abs/1503.04069
c. Understanding Deriving and Extending the LSTM: https://r2rt.com/written-memories-understanding-deriving-and-extending-the-lstm.html
d. Understanding LSTM Networks: http://colah.github.io/posts/2015-08-Understanding-LSTMs/
e. https://www.cse.iitm.ac.in/~miteshk/CS7015/Slides/Teaching/pdf/Lecture14.pdf
f. https://www.cse.iitm.ac.in/~miteshk/CS7015/Slides/Teaching/pdf/Lecture15.pdf

Module 5

a. GANs in Action: Deep Learning with Generative Adversarial Network Jakub Langgr, Vladimir Bok
b. Generative Deep Learning: Teaching Machines to Paint, Write, Compose, and Play David Foster
c. https://developers.google.com/machine-learning/gan

Course Contents and Lecture Schedule

No	Topic	No. of Lectures
1	**Module 1**	**8 Hours**
1.1	Review of Neural Networks: Model of a biological neuron	1
1.2	McCulloch Pitts Neuron, Activation functions	1
1.3	Perceptron, Perceptron Learning Algorithm	1
1.4	Convergence, Multilayer Perceptron	1
1.5	Back propagation	1
1.6	Learning XOR, Sigmoid Neurons	1
1.7	Gradient Descent, Feed forward Neural Networks	2
2	**Module 2**	**10 Hours**
2.1	Training Neural Networks	1
2.2	Initialization, Dropout	1
2.3	Batch normalization and drop out	1
2.4	Over fitting, under fitting, training and validation curves, data visualization, feature and weight visualization, tSNE	2
2.5	Introduction to TensorFlow, graphs, nodes, Tensor Data Structures - rank, shape, type	2
2.6	Building neural networks with tensor flow	2
2.7	Introduction to Keras	1
3	**Module 3**	**10 Hours**
3.1	Convolutional neural networks	1
3.2	Convolution operation	2
3.3	Back propagation in multilayer neural networks	1
3.4	Convolutional layers in neural network, pooling	2
3.5	Fully connected layers	2
3.6	Case study: Architecture of Lenet, Alexnet and VGG 16	2
4	**Module 4**	**8 Hours**
4.1	Recurrent neural networks	1
4.2	Back propagation: vanishing gradients, exploding gradients	1
4.3	Truncated Backpropagation Through Time	1
4.4	LSTM	1
4.5	Gated Recurrent Units (GRUs)	1
4.6	Long Short-Term Memory (LSTM) Cells	1
4.7	Solving the vanishing gradient problem with LSTMs	2
5	**Module 5**	**9 Hours**
5.1	Autoencoders, Variational autoencoders	2
5.2	Generative Adversarial Networks (GAN)	1
5.3	Discriminative and generative models	2
5.4	GAN Discriminator, GAN Generator, upsampling,	1
5.5	GAN Training	1
5.6	GAN challenges, Loss functions, cross entropy, minimax loss, Wasserstein loss	2