

K. J. Somaiya Institute of Technology, Sion, Mumbai-22
(Autonomous College Affiliated to University of Mumbai)

Nov – Dec 2025		
(B.Tech) Program: Artificial Intelligence and Data Science Scheme <i>Scheme: II / IIB</i>		
Examination: LY Semester: VII		
Course Code: AIC701 and Course Name: Deep Learning		
Date of Exam: 25/11/2025	Duration: 2.5 Hours	Max. Marks: 60

Instructions: (1) All questions are compulsory. (2) Draw neat diagrams wherever applicable. (3) Assume suitable data, if necessary.				
		Max. Marks	CO	BT level
Q 1	Solve any two questions out of three: (05 marks each)	10		
a)	Define deep learning. Explain different categories of Deep learning		CO1	U
b)	Explain the structure of a basic autoencoder. What are the encoder and decoder responsible for?		CO3	U
c)	What is a Recurrent Neural Network (RNN), and how is it designed to handle sequential data?		CO5	U
Q.2	Solve any two questions out of three: (05 marks each).	10		
a)	Explain the concept of backpropagation. How does it help in training MLPs?		CO1	U
b)	Apply max pooling and average pooling to the 7 x 7 image with stride=1 and padding=0. Filter will be of size 3 x 3. Compute output size of the image.		CO4	Ap
c)	Draw a RNN mathematical model for the word 'The India is my home country'		CO5	Ap
Q.3	Solve any two questions out of three. (10 marks each)	20		
a)	Compare SGD, mini-batch gradient descent, and full batch gradient descent in terms of computation, convergence, and stability.		CO2	U
b)	What is a contractive autoencoder, and how does it differ from standard and sparse autoencoders? How does a contractive autoencoder promote robustness to input variations?		CO3	U

K. J. Somaiya Institute of Technology, Sion, Mumbai-22
(Autonomous College Affiliated to University of Mumbai)

Nov – Dec 2025

(B.Tech) Program: Artificial Intelligence and Data Science Scheme

Examination: LY Semester: VII

Course Code: AIC701 and Course Name: Deep Learning

Date of Exam: 25/11/2025

Duration: 2.5 Hours

Max. Marks: 60

c)	What is padding in the context of convolutional layers, and why is it used? What is stride in the context of convolution, and how does it influence the output dimensions?		CO4	U
Q.4	Solve any two questions out of three. (10 marks each)			
a)	Compare L1 and L2 regularization. When might you choose one over the other? How do L1 and L2 regularization reduce the risk of overfitting?	20	CO2	U
b)	What is Backpropagation Through Time (BTT), and how is it used to train RNNs? Describe the process of applying BTT to an unfolded RNN over multiple time steps.		CO5	U
c)	What are deep fakes, and how are they generated using GANs? Discuss the technological advancements that have enabled the creation of realistic deep fakes.		CO6	U
