



**Objectives**

- To get insights of image and video analysis task using deep learning.
- To get the students aware about building blocks used in deep learning based solutions.

Unit No.	Topic(s)	No. of Hours
1.	<b>Introduction</b> What is Deep Learning?, Why Deep Learning?, What is a neural network?, Deep Learning Success Stories, McCulloch Pitts Neuron, Thresholding Logic, Perceptron, Perceptron Learning Algorithm	4
2.	<b>Deep Feedforward Networks and Regularization in Deep Learning</b> Example: Learning XOR, Gradient-based learning, Hidden units, Backpropagation and other differentiation algorithms, Parameter Norm Penalties, Norm Penalties as Constrained Optimization, Regularization and Under-Constrained Problems, Dataset Augmentation, Noise Robustness, Semi-Supervised Learning, Multitask Learning, Early Stopping, Parameter Tying and Parameter Sharing, Bagging and Other Ensemble Methods, Adversarial Training	10
3.	<b>Convolutional Networks</b> The Convolution Operation, Pooling, Convolution and Pooling as an Infinitely Strong Prior, Variants of the Basic Convolution Function, Structured Outputs, Data Types, Efficient Convolution Algorithms, Random or Unsupervised Features, The Neuroscientific Basis for Convolutional Networks	8
4.	<b>Sequence Modelling: Recurrent and Recursive Nets</b> Unfolding Computational Graphs, Recurrent Neural Networks, Bidirectional RNNs, Encoder-Decoder Sequence-to-Sequence Architectures, Deep Recurrent Networks, Recursive Neural Networks, The Challenge of Long-Term Dependencies, Echo State Networks, Leaky Units and Other Strategies for Multiple Time Scales, The Long Short-Term Memory and Other Gated RNNs, Optimization for Long-Term Dependencies, Explicit Memory	12
4.	<b>Deep Learning Applications, Platforms and Software Libraries</b> Large-scale deep learning, Computer vision, Speech recognition, Natural language processing, H2O.ai, Dato GraphLab, Theano, Caffe	6



### Reference Books

1. Deep Learning  
by Ian Goodfellow, Yoshua Bengio and Aaron  
An MIT Press book
2. Neural Networks and Deep Learning  
by Michael Nielsen  
<http://neuralnetworksanddeeplearning.com>
3. Pattern Classification  
by Richard O. Duda, Peter E. Hart, David G. Stork  
John Wiley & Sons Inc

### Outcomes

After completion of subject, students would be able to:

- understand and apply multilayer perceptron, convolutional neural networks, recurrent neural networks.
- acquire the knowledge of applying deep learning techniques to solve various real life problems.
- use the functionalities provided by deep learning platforms and software libraries.

**Suggested list of Practical (at least 10 practical are to be performed by students. These practical should cover majority of all topics of syllabus.)**

**This is the suggested list of practical but it may not be limited only to this list.**

1. Implement the sigmoid function using numpy. Implement the function sigmoid\_grad() to compute the gradient of the sigmoid function with respect to its input.
2. Perform following operations on two vectors and compare the computation time.
  - (i) Classic dot product of vectors
  - (ii) Classic outer product
  - (iii) Classic elementwise multiplication
  - (iv) Classic general dot product
3. Implement the numpy vectorised version of the L1 and L2 loss.
4. Write a program to generate XOR function using McCulloch-Pitts neuron and appropriate values for weights, bias and threshold.
5. Write a program to recognize the number from 0; 1; 2; 3; : : : ; 9. A number is represented as a  $5 \times 3$  matrix of 0 and 1. For any valid point it is taken as 1 and



invalid point it is taken as 0. The net has to be trained to recognize all the numbers and when the test data is given. The file has three components, as:

- input-data: The training data. Each column represents a number.
- output-data: A  $10 \times 10$  matrix of desired outputs.
- test-data: Test data. Each column is a test vector.

6. For given a dataset containing:

- a training set of images labeled as cat ( $y=1$ ) or non-cat ( $y=0$ )
- a test set of images labeled as cat or non-cat
- each image is of shape (num\_px, num\_px, 3) where 3 is for the 3 channels (RGB).

Thus, each image is square (height = num\_px) and (width = num\_px).

You have to build a simple image-recognition algorithm (logistic regression) that can correctly classify pictures as cat or non-cat.

7. Write a program to build a deep neural network, with as many layers as you want. Use non-linear units like ReLU to improve your model.
8. Write a program to build a deep network, and apply it to cat vs non-cat classification. Compare the accuracy of DNN with logistic regression of practical 6.
9. Implement convolutional (CONV) and pooling (POOL) layers in numpy, including both forward propagation and backward propagation.