Deep learning - IIT Ropar

PROF. SUDARSHAN IYENGAR
Department of Computer Science and Engineering
IIT Ropar

TYPE OF COURSE : Rerun I Elective I UG/PG
COURSE DURATION : 12 weeks (26 July' 21 - 15 Oct' 21)
EXAM DATE : 23 Oct 2021

PRE-REQUISITES : Working knowledge of Linear Algebra, Probability Theory. It would be beneficial if the participants have done a course on Machine Learning.

INTENDED AUDIENCE : Any Interested Learners

COURSE OUTLINE :
Deep Learning has received a lot of attention over the past few years and has been employed successfully by companies like Google, Microsoft, IBM, Facebook, Twitter etc. to solve a wide range of problems in Computer Vision and Natural Language Processing. In this course we will learn about the building blocks used in these Deep Learning based solutions. Specifically, we will learn about feedforward neural networks, convolutional neural networks, recurrent neural networks and attention mechanisms. We will also look at various optimization algorithms such as Gradient Descent, Nesterov Accelerated Gradient Descent, Adam, AdaGrad and RMSProp which are used for training such deep neural networks. At the end of this course students would have knowledge of deep architectures used for solving various Vision and NLP tasks

ABOUT INSTRUCTOR :
Sudarshan Iyengar has a PhD from the Indian Institute of Science and is currently working as an Assistant Professor at IIT Ropar and has been teaching this course from the past 4 years.

'U 3DGPDYDWL UHFHLYHG % (GHJUHH LQ &RPSXWHU 6FLHQFH DPS
*XOEDUJD 9 7 8 %HOJDXP .DUQDWDND 0 (GHJUHH LQ &RPSXWHU 6
3DQMDE (QJLQHHULQJ &ROOHJH 'HHPHG WR EH 8QLYHUVLW\ &KDQG

COURSE PLAN :

Week 1: History of Deep Learning, Deep Learning Success Stories, McCulloch Pitts Neuron

Week 2: Multilayer Perceptrons (MLPs), Representation Power of MLPs, Sigmoid Neurons, Gradient Descent

Week 3: Feed Forward Neural Networks, Back propagation

Week 4: Gradient Descent (GD), Momentum Based GD, Nesterov Accelerated GD, Stochastic GD

Week 5: Principal Component Analysis and its interpretations, Singular Value Decomposition

Week 6: Auto encoders and relation to PCA, Regularization in auto encoders, Denoising auto encoders, Sparse auto encoders

Week 7: Regularization: Bias Variance Tradeoff, L2 regularization, Early stopping, Dataset augmentation

Week 8: Greedy Layerwise Pre-training, Better activation functions, Better weight initialization methods, Batch Normalization

Week 9: Learning Vectorial Representations Of Words

Week 10: Convolutional Neural Networks, LeNet, AlexNet, ZF-Net, VGGNet, GoogLeNet, ResNet

Week 11: Recurrent Neural Networks, Back propagation through time (BPTT), Vanishing and Exploding Gradients, Truncated BPTT, GRU, LSTMs

Week 12: Encoder Decoder Models, Attention Mechanism, Attention over images