Lahore University of Management Sciences

CS 437/5317– Deep Learning

Fall 2018-19

Instructor	Murtaza Taj
Room No.	9-G11A
Office Hours	TBA
Email	murtaza.taj@lums.edu.pk
Telephone	3301
Secretary/TA	Danish Farid (19100172), Rabeez Riaz (19100165)
TA Office Hours	TBA
Course URL (if any)	LMS

Course Basics				
Credit Hours	3			
Lecture(s)	Nbr of Lec(s) Per Week	2 (MonWed)	Duration	75 min (4:00 – 5:15 pm)
Recitation/Lab (per week)	Nbr of Lec(s) Per Week	0	Duration	N/A
Tutorial (per week)	Nbr of Lec(s) Per Week	0	Duration	N/A

Course Distribution	
Core	No
Elective	Yes
Open for Student Category	Graduate, CS 200, Math 120
Close for Student Category	Freshman

COURSE DESCRIPTION

Deep Learning is a hierarchical learning methodology based on artificial neural networks which are algorithms inspired by the structure and function of the brain. It has applications in wide-range of industries these days such as face-recognisers working at massive scales, robotics, speech translation, text analysis, improving customer experience, autonomous vehicles etc.

In this course we will take a "hands-on approach" and start will implementation of basic building blocks such as training a simple perceptron and move to design and train a deep convolution neural network. Course will concentrate in developing both mathematical knowledge and implementation capabilities. The implementations will be python based using TensorFlow and Keras. After establishing our foundation in convolutional neural networks we will start looking into applications of deep learning in both spatial as well as time-series data and explore various network architectures suited for each. The objective is to help you build a career in AI and Machine learning, to make you comfortable enough that you can understand various learning problems and develop your own deep learning based solutions.

COURSE PREREQUISITE(S)

- CS 200 - Introduction to Programming
- Math 120 – Linear Algebra with Differential Equations

COURSE OBJECTIVES

- Motivate the class about data driven problem solving paradigm
- Introduce the basic theory and applications of Deep Learning
- Provide a solid foundation for further work in this area both in academia and industry

Learning Outcomes

- To familiarize students with the concepts of Convolutional Neural Network and the processing going on within each layer
- Make students confident that they can solve machine learning problems, through the use of deep features

Lahore University of Management Sciences

Grading Breakup and Policy	
Assignment(s):	40%
HomeWork:	05%
Quiz(s):	05%
Midterm Examination:	20%
Final Examination:	25%

Examination Detail	
Midterm Exam	Yes/No: Yes Combine Separate: N/A Duration: 120 min Preferred Date: 7th Week Exam Specifications: Written
Final Exam	Yes/No: Yes Combine Separate: Comprehensive Duration: 180 min Exam Specifications: Written

Lahore University of Management Sciences

S#	Module	L#	Topics Covered (tentative)	Book Chapters/ Readings
	Introduction	1	Overview – Brain – Neuron - Hubel & Wiesel, 1959 What is Learning? What is Machine Learning Historical Context Shallow Feature Learning What is Classification? Shallow vs. Hierarchical vs. Deep Features (ML vs. DL) Course Logistics Success Stories	[Buduma] Ch 1
			Tutorial: Pandas, Jyupter Notebook Python/Numpy	
2	Neural Networks	1	What is Regression – Line fitting? The Neuron – Biologically Perceptron Linear Perceptron as Neuron Logistic Regression	[Buduma] Ch 1
	Training Feed-Forward NN	1	The Fast-Food Problem (Hinton / Buduma) Gradient Descent - Intuition The Delta Rules and Learning Rules Handout: MSE with Sigmoid, cross entropy+sig, MSE+Softmax, cross entropy+softmax Gradient Descent with Sigmoid Neurons More Derivative Examples (Ng) Computation Graph Derivatives with a Computation Graph Multi-layer Perceptron Gradient Descent The Back Propagation Algorithm Stochastic and Minibatch Gradient Descent	[Buduma] Ch 2
			Tutorial: Computing a Neural Network's Output	
			PA1: Implementing Feed-Forward Neural Network with Back Propagation in Python	
3	Practical Aspects	2	Test Set, Validation Set, Overfitting Regularization Hyper parameter tuning Data Augmentation Vanishing/Exploding Gradients Weight/Initialization Methods???? Activation Functions Softmax Optimization Algos Gradient Descent with momentum Learning rate adaptation (AdaGrad, RMSProp, Adam)	[Buduma] Ch 2 [Buduma] Ch 4
			Tutorial: TensorFlow, Implementing NN in TensorFlow+Keras	[Buduma] Ch 3
			PA2: Neural Network and its hyper parameter tuning	

Lahore University of Management Sciences

S#	Module	L#	Topics Covered (tentative)	Book Chapters/ Readings
4	ConvNet	3	Convolution-1D Convolution-2D Convolution-Filters (Edge detection) Forward and Backward Propagation using Convolution operation Transforming Multilayer Perceptron to Convolutional Neural Network Texture Classification Example + Filter Banks (Dr. Mohsen) A toy ConvNet: X's and O's Eg. (Brandon Rohrer)/ Full Arch Description on ConvNet Feature Maps Pooling, FC, Batch Normalization etc Closing the loop on MNIST with ConvNet Accelerating training with batch normalization Multi-Class Learning - Building a ConvNet for CIFAR-10 Transfer Learning	[Buduma] Ch 5
			PA3: Implement 2D Conv, with MNIST+CIFAR 10, with batch generator	
5	ConvNet: Case Studies	2	Classical CNN: Case Studies AlexNet, VGG, GoogleNet, ResNet, ResNet, Inception, U-Net Relationship between ConvFilters and Receptive Field	
			Tutorial: Visulaization (T-SNE, etc)	
			PA4: Image Classification in Remote Sensing Imagery	
6	Object Detection	2	Classification + Regression Region Proposal Networks RCNN Faster RCNN YOLO: You Only Look Once SSD: Single Shot Detector	
			PA5: Object Detection in Remote Sensing Imagery	
7	Autoencoders	3	Embedding and Representation Learning Learning Lower-Dimensional Representations 117 Principal Component Analysis 118 Motivating the Autoencoder Architecture 120 Implementing an Autoencoder in TensorFlow 121 Denoising to Force Robust Representations 134 Sparsity in Autoencoders Stacked Autoencoders Image Segmentation Instance Segmentation Image Retrieval using Unsupervised/Semi-supervised Learning	[Buduma] Ch 6
			Tutorial: Google Earth Engine	
			PA6: Image Segmentation AND/OR Image Retrieval	

Lahore University of Management Sciences

S#	Module	L#	Topics Covered	Book Chapters/ Readings
8	GANS	2	GANS Adversarial Learning Adversarial Attacks Adversarial AAE Discriminative AAE Cyclic GANS – CT+X-Ray DC GANS Cross-View Generation – Remote Sensing	
9	Sequence Models	4	Intro to RNN & Back Prop through time Character RNN Difference Units Recurrent Neural Networks 173 The Challenges with Vanishing Gradients 176 Long Short-Term Memory (LSTM) Units 178 TensorFlow Primitives for RNN Models 183 Implementing a Sentiment Analysis Model 185 Solving seq2seq Tasks with Recurrent Neural Networks 189 Augmenting Recurrent Networks with Attention 191 Dissecting a Neural Translation Network Deep Captioning Visual Question Generation	[Buduma] Ch 7
			PA7: Natural Language Processing using LSTM	
10	Memory Augmented NN	3	Neural Turing Machines 219 Attention-Based Memory Access 221 NTM Memory Addressing Mechanisms 223 Differentiable Neural Computers 226 Interference-Free Writing in DNCs 229 DNC Memory Reuse 230 Temporal Linking of DNC Writes 231 Understanding the DNC Read Head	[Buduma] Ch 8
11	Deep Reinforcement Learning	2	[Tentative] Q-Learning Tool box - PAC MAN	[Buduma] Ch 9

Lahore University of Management Sciences

Textbook(s)/Supplementary Readings

***Required**:* [Buduma] Fundamentals of Deep Learning: Designing Next-Generation Machine Intelligence Algorithms 1st Edition by Nikhil Buduma

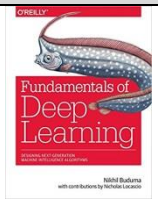

***Reference**:*

1. Neural Networks and Deep Learning by Michael Nielsen
http://neuralnetworksanddeeplearning.com/

2. Deep Learning, An MIT Press book, by Ian Goodfellow and Yoshua Bengio and Aaron Courville
http://www.deeplearningbook.org/