Deep Learning

Instructor: Cesar Acosta-Mejia

Course Description

This course introduces Deep Learning (DL) fundamentals and applications. Students will learn the elements of neural networks (NN), the methods used to optimize their performance, the most common NN architectures, and applications of DL models to analytics and AI. These applications range from regression, classification, decision making, image classification, and image identification. Students will learn how to implement these applications using python library Keras.

Learning Objectives and Outcomes

At the end of the course the students will know

- the fundamental principles of NN models
- the elements of a neural network model
- the most common types of NN architectures and what type of problems each is suitable for
- how to build and implement NN models for prediction and classification business data analytics problems
- how to build and implement NN models for artificial intelligence applications

Prerequisite(s): None

Recommended Preparation: Undergraduate courses as follows

- It is strongly recommended to have python programming experience
- Calculus to the level of MATH 229 or equivalent
- Statistics to the level of ISE 225 or equivalent

Course Notes

The course material is available on Blackboard.

Technological Proficiency and Hardware/Software Required

Required software: Python programming language is used throughout the course.
Jupyter Notebook is used as the main interface for documenting the scripts and results.

Textbook (required)

- Chollet F., *Deep Learning with Python,* 2nd Ed., Manning, 2021, ISBN 978-16172996864

Supplementary Materials (for reference only)

- Goodfellow I. et al., *Deep Learning,* MIT Press, 2017, ISBN 978-0262035613
- VanderPlas, *Python Data Science Handbook*, O'Reilly, 2017

Description and Assessment of Assignments

All assignments and examinations are on-line. Unless otherwise noted the assignments are individual. All homework assignments are released and distributed one-week in advanced. Dates are shown in the Course schedule on page 4. Submit on to Blackboard by the due date. No late homework will be accepted.

Grading Policy

Assignment	Points	% of Grade
Homework	100 each (6^+ homework assignments)	30
Midterm	100	30
Final	100	40
TOTAL		100

Grading Scale (Course final grades will be determined using the following scale)

A	94-100	B-	80-82.9	D+	67-69
A-	90-93.9	C+	77-79	D	63-66.9
B+	87-89	C	73-76.9	D-	60-62.9
B	83-86.9	C-	70-72.9	F	59.9 and below

Assignment Submission Policy

Assignments should be typewritten and clean. Email submissions and late submissions are not allowed. No make-up exams are considered.

Timeline and Rules for submission

Assignments must be submitted on the due date. Solutions will be released soon after the homework due date.

Course Schedule: A Weekly Breakdown

	Date	Topics/Daily Activities	Deliverables	slides	ref. and files
1	**June 28**	**Introduction.** Quick review of Python and Jupyter Notebook (JN). Library numpy. Gradient Descent.	HW1 Python and numpy	overview.ppt python.ppt numpy.ppt python basics.pdf	intro.ipynb Ch 3.4 numpy.ipynb numpyreg.ipynb Tour-on-python.pdf
2	**July 3** **(**recorded**)**	**Generalization**. Overfitting. Regularization. Holdout cross validation. Pandas review.		generalization.ppt rr.ppt pandas basics.pdf	Ch 5.1, 5.2, 5.4 ridge9.ipynb cancerRR.ipynb hitters.csv, small.csv
3	**July 5**	**Intro to Neural Networks (NN)** The Densely connected NN. The NN as a NL regression model. The Gradient. Building a NN with sklearn.	HW1 due	nlreg.ppt nn4.ppt	Kutner Ch. 13 Ch. 2.2, 2.3,2.4 gradient.ipynb perceptron.ipynb multilayerp.ipynb
4	**July 10**	**Introduction Deep Learning (DL).** The MLP. A first DL model with Keras. Multilabel classification models.		mnist.ppt	Ch 2.1, 2.4, 2.5 mnist.ipynb
5	**July 12**	**Deep Learning applications.** NN for classification. Binary and Multiclass classification models.	HW2	classification.ppt	Ch 4.1, 4.2 imdb.Ipynb reuters.ipynb
6	**July 17**	**Deep Learning applications.** NN for Regression problems. K-fold cross validation.		dlreg.ppt cv3.ppt	Ch 4.3, 5.2 boston.ipynb
7	**July 19**	**Convolutional Neural Networks (CNN).** Basic Components of a CNN. CNN applications. Image classification.	HW2 due Midterm release	cnn.ppt	Ch. 8. vision.ipynb Ch 9. image.ipynb
8	**July 24**	**Deep learning for Time Series.** Forecasting application. Introduction to Recurrent NN.		lstm.ppt	Ch. 10 temperature. ipynb
9	**July 26**	**Deep learning for text data.** Natural language processing. Pre-procesing text data.	Midterm due	text.ppt	Ch. 11.2 nlp.ipynb
10	**July 31**	**Deep learning for sequences.** Processing a sentence as a sequence. Sequence Models.	HW3	seq.ppt	Ch. 11.3 sequences.ipynb
11	**Aug 2**	**Transformers.** Transformer encoders. Sequence to sequence learning	HW3 due Final release.	transf.ppt	Ch. 11.4 transformer.ipynb
12	**Aug 7**	**Generative Deep learning.** Generative Adversarial Networks.		gdl.ppt	Ch. 12
13	**Aug 9**	**FINAL EXAM**			