

**COLUMBIA BUSINESS SCHOOL**  
**ARTIFICIAL INTELLIGENCE - B9654 - Spring 2024**

**Dr. George A. Lentzas**

**Email: [gl2543@columbia.edu](mailto:gl2543@columbia.edu)**

**Course Description:** This course is the second of two courses that will introduce business students to the exciting and growing literature in Machine Learning / Artificial Intelligence, with a focus on applications in finance, marketing and business. We will cover advanced topics in neural networks, deep learning and artificial intelligence and their applications in finance, marketing and business.

**Course Requirements:** Students are expected to have completed the Spring Semester A-Term Machine Learning course offered by the instructor (similar graduate level courses will be acceptable with instructor approval). Students should be familiar with probability theory, matrix algebra, and basic concepts/models in machine learning.

**Required Text:** The main readings for this course reading are the (2<sup>nd</sup> edition) textbooks "Deep Learning with R" and "Deep Learning with Python" by F. Chollet and J.J. Allaire and F.Chollet respectively, referred to **DLRP**. Additional will be the textbook "Deep Learning" by Goodfellow, Bengio and Courville, referred to as **DL**, (<http://www.deeplearningbook.org/>).

**Class Preparation:** Students are expected to have completed the session reading ahead of class to facilitate class participation and discussion. You need to take this seriously to be able to make the most out of this class as we will cover a lot of very advanced material.

**Research Paper:** Students will write a research paper that addresses a well-defined finance or marketing question using the deep learning /

artificial intelligence toolkit. The Research Paper grade will count as per the formula below.

**Class Quizzes**: At the end of each class there will be a 5 minute multiple choice quiz that will count towards your grade.

**Grade** =  $1/4 * \text{Class Quizzes} + 1/4 * \text{Final} + 1/4 * \text{average of [HW1, HW2]} + 1/4 * \text{Research Paper}$

Homework and Research Paper is due promptly. **Late submissions will be penalized by 5% per 30-minute delay.**

**Office Hours**: By Appointment

## **SCHEDULE**

### **SESSION 1 (Neural Networks: Theory)**

- ☐ An introduction to AI & Neural Networks
- ☐ Deep Neural Networks & Deep Learning
- ☐ Regularization and the Bias-Variance Tradeoff
- ☐ Reading: Elements of Statistical Learning Ch. 11 | DL Ch. 1-3, 4.3-4.5, 5, 6

### **SESSION 2 (Neural Networks: Practice)**

- ☐ Introduction to Tensors
- ☐ A Closer Look to Optimization for Deep Learning
- ☐ Implementing Deep Neural Networks in Keras
- ☐ Reading: DL Ch. 7-8 | DLRP Ch. 1-5

### **SESSION 3 (Convolutional Networks)**

- ☐ Introduction to Convolutional Networks
- ☐ Lab: Implementing a simple Convolutional Network in Keras
- ☐ Reading: DL Ch. 8-9 | DLRP Ch. 8-9

#### **SESSION 4 (Recurrent Networks Part I)**

- ☐ Introduction to Recurrent Networks
- ☐ Long Short Term Memory Models and Gated RNNs
- ☐ Deep Learning Practical Issues and Applications
- ☐ Lab: Implementing an LSTM in Keras
- ☐ Reading: DLRP Ch. 10-11

#### **SESSION 5 (Recurrent Networks Part II)**

- ☐ The Transformer and Attention Architectures
- ☐ Lab: Implementing a Transformer in Keras
- ☐ Reading: DLRP Ch. 11-12

#### **SESSION 6 (Large Language Models)**

- ☐ History of Transformer Models
- ☐ Introduction to LLMs
- ☐ Chat GPT and the rest
- ☐ Generative Learning
- ☐ Reading: Lecture Notes