IFI 8420 – Machine Learning and Deep Learning for Business
Course Syllabus – Spring 2021

(Draft – Subject to Change)

Instructor:	TBD
Class Schedule:	TBD
Classroom:	TBD
Office Hours:	TBD
Course Description:	This course provides an introduction to machine learning and artificial intelligence to help you understand the implications of these new technologies for business strategy and deploying these methods to create value. The topics covered include supervised and unsupervised learning, linear and nonlinear regression/classification and boosting techniques. The methods are taught through business applications including time series, stock market value/direction prediction, real-estate price prediction, fraud detection and user signup in complex systems. Since unstructured data comprises a large and important source of information, we cover deep learning techniques to tackle complex problems such as image understanding and analytics, text analytics, and recommendation systems. The unique set of problem solving skills introduced in this course will help the students derive novel insights from their business problems.
Course Objectives:	After taking this course, the students are expected to: • Have a good understanding of the basics of machine learning and its application to business problems and business data sets • Have the insight to become involved in more complex AI projects • Understand the pros and cons of different techniques and their preferred business application scenarios
Contributing Texts:	James, G., Witten, D., Hastie, T., Tibshirani, R., \An Introduction to Statistical Learning: with Applications in R", Springer, 2013 (main text). URL (E-book): https://goo.gl/8NYEo4 (Deep Learning) Goodfellow, Ian, Bengio, Yoshua, and Courville, Aaron, \Deep learning", MIT press, 2016. URL (E-book): https://www.deeplearningbook.org

(Further Reading) Kelleher, John D., Mac Namee, Brian, and D'arcy, Aoife, \Fundamentals of machine learning for predictive data analytics: algorithms, worked examples, and case studies". MIT press, 2015.

Class Schedule and Activities:

CLASS	TOPIC AND COURSE CONTENT
Session 1	Introduction to machine learning and foundations: why machine learning in business?
Session 2	Linear regression and predicting real estate price
Session 3	More on Linear regression and extensions, time series analysis, stock market value prediction
Session 4	Classification (logistic regression, KNN, LDA, QDA): fraud detection and stock market direction prediction
Session 5	More on classification, support vector machine: fraud detection, customer churn prediction
Session 6	Bootstrap, Cross Validation and model selection to identify the key business factors
Session 7	Handling nonlinear business models with few input features, beyond linearity (polynomials and splines)
Session 8	Overview of the material
Session 9	Midterm Exam
Session 10	Semester Break
Session 11	Interpretability of business models: random forests and decision trees, predicting user signups
Session 12	Deep learning and image analytics, object detection, recommendation systems
Session 13	More cutting-edge business solutions including patient diagnosis, text analytics; etc.
Session 14	More on neural networks, deep learning and variants, intro to GANs and learning the distribution of business interactions
Session 15	Unsupervised learning (PCA, K-means, ...) for market segmentation
Session 16	Class summary + Final presentations