Course Title:
Introduction to Machine Learning in Python (MBA)

Course Number: 099102
Number of credit points: 2
Mini-Semester: of the Academic Year: **2019-2020**
Time: 14:00-17:45

Course Instructor: Omer Ben-Porat
E –mail: **omerbp@campus.technion.ac.il**

Meeting time for students: right after class

Syllabus:

Course Objectives: This course introduces students to advanced Machine Learning techniques in Python. Our focus will be on classification algorithms (SVM, Logistic Regression), regression algorithms (Linear Regression, Lasso, Ridge), and ensemble methods (Random Forest). Moreover, we will explore the software tools that constitute the Python data science ecosystem.

Course Content & Scope: As the course is highly applicative, potential candidates must be proficient in at least on programming language (preferably dynamic one). Advanced mathematical background (linear algebra, probability) is also assumed.

Learning Objectives: By the end of the course the student will know how to:
1. Write basic code in Python
2. Use standard data science packages in Python, e.g., Pandas, Numpy and Scikit-learn
3. Conduct exploratory data analysis
4. Employ predictive algorithms for regression and classification
5. Understand and analyze the challenges in model evaluation

Teaching Methods: Presentations, hands-on sessions in the classroom.

Teaching Materials: Slides, Jupiter notebooks (Python).

Readings (Recommended):
- https://stackoverflow.com
- Python data science handbook: essential tools for working with data (Jake VanderPlas).

Student Assessment:

Your grade will be composed of classwork grade (10%), proficiency exam (50%), and individual course project (40%).

- The exam will encompass questions on Python's syntax, implementation of simple algorithms (e.g., matrix multiplication), and foundations of Machine Learning techniques.
- In the Analytics project, you will apply algorithms learned in class to a dataset of your choice (preferably related to your domain of expertise). Your goal is to explore the data to gain insights, and develop a predictive model for a problem or question that arises in context of your data. More information will be given on the first week.

Course Plan

1. Introduction to Python programming (overview, basic syntax, variable types, containers, objects, data structures, Jupyter).

2. Exploratory Data Analysis (EDA) (input/output, Numpy, Pandas, Matplotlib).

3. Supervised learning (Learning Framework, Linear Regression, Logistic Regression, Model Selection with Cross-validation, Decision Trees, Bootstrapping, Random Forests, Gradient Boosting)

4. If time permits, one or more of the following subjects: unsupervised learning (K-means, Hierarchical methods), text analysis (sentiment analysis), Reinforcement Learning (Multi-Armed Bandits), ethics in Machine Learning.