Course Title	Principles of Machine Learning				
Course Code	AEM675				
Course Type	Elective				
Level	Master (2nd cycle)				
Year / Semester	2nd Year / 3rd Semester				
Teacher's Name	TBA				
ECTS	10	Lectures / week	3 Hours / 14 weeks	Laboratories / week	None
Course Purpose and Objectives	The course introduces the fundamental concepts, theory, and algorithmic ideas of Machine Learning. It provides the student both with a foundation for either applying ML techniques on real-world problem or performing research on developing novel ML algorithms. It also forms a foundation for several other Data Science and AI courses, including advanced topics in ML, natural language processing, Big Data Analytics and others. Specifically, the course focuses on supervised classification techniques, basic and advanced classifiers (logistic regression, naïve Bayes classifier, K-nearest neighbors, support vector machines, decision trees, random forests), statistical hypothesis testing, metrics of performance (ROC curves and AUC), estimation of performance and tuning of hyper-parameters (cross-validation, nested cross validation, and bootstrap bias corrected CV), and feature selection (forward-backward search, lasso, orthogonal matching pursuit).				
Learning Outcomes	Upon successful completion of this course students should be able to: ● Define basic ML tasks and types of analysis, such as supervised learning, unsupervised learning, reinforcement learning, classification and regression, and feature selection. ● Discuss the inner workings of standard ML classification and feature selection algorithms. ● Illustrate how to solve the problem of selecting algorithms, tuning their hyper-parameters, and estimating the performance of the final predictive model. ● Perform and apply ML pipelines to real-world problems, dealing with problems such as representing the problems as an ML task, representing appropriately the data, applying and tuning an ML pipeline, and interpreting results. ● Define key statistical estimation and hypothesis testing concepts, with a focus on the ones that are routinely employed within ML algorithms. ● Have a solid, foundational basis to perform ML research and proceed with other courses that employ ML algorithms and concepts.				
Prerequisites	None		Co-requisites	None	
Course Content	1) Introduction to ML, supervised, unsupervised, reinforcement learning, hypothesis (models) spaces, examples of ML applications				

	2) Probability theory and concepts for ML, axioms of probability, conditional probability, Bayes theorem, maximum likelihood estimation, maximum a posteriori estimation 3) Logistic Regression and fitting with gradient descent 4) Hypothesis testing, and permutation-based hypothesis testing 5) Naïve Bayes and K-Nearest Neighbors 6) Decision Trees and Random Forests 7) Metrics of performance, Receiver Operating Characteristic Curves (ROC), and Area Under the ROC curve 8 and 9) Estimation of performance and hyper-parameter tuning using cross validation techniques 10) Basics of optimization and constrained optimization 11, 12) Support Vector Machines 13) Basic Feature Selection All lectures will consist of a theoretical part presenting concepts and techniques and a practical part were the ML techniques will be applied for problem solving.
Teaching Methodology	Face to Face
Bibliography	Machine Learning, Tom Mitchell, McGraw Hill, Latest Edition Pattern Recognition and Machine Learning, Christopher Bishop, Springer, Latest Edition The Elements of Statistical Learning, Jerome H. Friedman, Robert Tibshirani, and Trevor Hastie, Latest Edition, Springer An Introduction to Statistical Learning, with Applications in R. Gareth James, Daniela Witten, Trevor Hastie and Robert Tibshirani, Springer, Latest Edition
Assessment	Examinations 50% Assignments 40% Class Participation and Attendance 10% 100%
Language	English