Course Title	Artificial Neural Networks and Deep Learning				
Course Code	AI625				
Course Type	Elective				
Level	Master (2nd Cycle)				
Year / Semester	1st Year/2nd Semester or 2nd Year/1st Semester				
Teacher's Name	TBA				
ECTS	7	Lectures / week	Up to 6 Teleconferences	Laboratories / week	None
Course Purpose and Objectives	This course has a two-fold purpose: first it presents the fundamental aspects of Artificial Neural Networks (ANN) and second it introduces to more advanced topics such as Deep Learning (DP) Networks. After a short review of conventional neural networks and learning processes, the course will introduce modern practices for deep networks, including training, optimization, convolutional networks, recurrent and recursive nets. Furthermore, the course will focus on practical methodologies concerning the design, data preprocessing, hyperparameter selection, implementation and performance evaluation of a Deep Learning system as well as applications of deep learning techniques to real world problems, such big data mining, image processing and natural language processing.				
Learning Outcomes	Upon succesful completion of this course students should be able to: • Recall the basic algorithms and methods of Artificial Neural Networks basics as well as their training algorithms. • Discuss, explain and report various deep learning algorithms for a specific problem. • Choose and interpret a suitable algorithm/method in order to meet problem's specifications. • Analyze and combine known technics in order to face real world problems. • Analyze and preprocessing the given data to fit in ANN and DL algorithms/methods. • Put the different system parts (preprocessed data, implemented, algorithms, user interface) together in order to create a new operational learning system. • Evaluate the performance of the developed learning system. • Discuss the fundamental concepts of several types of deep neural networks. • Apply Deep Learning approaches to a variety of tasks.				
Prerequisites	AI605		Co-requisites		

Course Content	**Machine Learning Basics:** Presentation of some perspectives on traditional machine learning techniques, such as Neural Networks (NNs) that have strongly influenced the development of deep learning algorithms. After a short introduction to NNs, we will discuss the model of a neuron, and network architectures. Then, the different types of learning processes are presented. Finally, the several aspects concerning the training of a single layer perceptron are discussed. **Deep Feedforward Networks:** Presentation of deep learning neural network modes for function approximation. A simple learning example and the gradient-based learning are discussed, as well as other aspects like hidden units and architecture design. Then follows the foundation of the Back Propagation algorithm for deep learning and its variations. The relative algorithms are analyzed extensively, and implementation aspects are discussed. **Regularization for Deep Learning:** Presentation of selected advanced techniques for regularization and optimization of deep network models, such as parameter's norm penalization, norm penalties as constrained optimization and dataset augmentation. Furthermore, the semi-supervised learning paradigm and feature extraction techniques, bagging and ensemble methods are discussed. **Optimization for Training Deep Models:** Several challenges of training optimization, such as parameter optimization and adaptive learning rates will be discussed. Furthermore, the relative algorithms are presented and analyzed as well as optimization strategies and meta-algorithms. **Convolutional Neural Networks (CNNs):** Introduction to convolutional networks for scaling to large data sets. Presentation of the main building blocks of CNNs such as convolutional filters and their characteristics (stride, depth, width), activation functions, pooling operator. Several aspects of convolution operation are discussed and efficient algorithms for random or unsupervised features are presented as well as the neuroscientific basis of convolution networks. **Sequence Modeling: Recurrent and Recursive Neural Nets (RNNs):** Deep recurrent and recursive neural networks for temporal sequences processing are presented. The challenge of long-term dependencies, the description of long sort-term memory and other gating mechanisms as well as optimization aspects are discussed. **Practical Methodology:** General guidelines for the practical methodology involved in designing, building and configuring an application involving deep learning are discussed. These aspects include performance metrics, baseline models, gathering more data, hyperparameters' selection and debugging strategies. An example shows how these aspects are faced.

	Applications: A review of deep learning applications, such as Large-scale deep learning, computer vision, speech recognition. Natural language processing and other applications are presented.
Teaching Methodology	E-Learning
Bibliography	"Neural Networks: A Comprehensive Foundation", by Simon Haykin (Pearson Prentice Hall & online at: https://www.pdfdrive.com/neural-networks-a-comprehensive-foundationpdf-e18774300.html). "Deep Learning", by Ian Goodfellow, Yoshua Bengio and Aaron Courville (The MIT Press & online at http://www.deeplearningbook.org/). "Neural Networks and Deep Learning", by Michael Nielsen (online at http://neuralnetworksanddeeplearning.com/). "Deep Learning with Python" by Francois Chollet (Manning Publications).
Assessment	Final Examination 50% Assignments/On-going evaluation 50% 100%
Language	English