



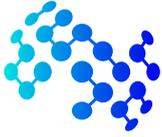
Core Courses Syllabi

AI702 - Deep Learning

Title	Deep Learning
Code	AI702
Loading	4 Credit-hours
Prerequisites	<ul style="list-style-type: none">Basics of Linear Algebra, Calculus, Probability and StatisticsProficiency in Python
Catalog Description	This course provides a comprehensive overview of different concepts and methods related to deep learning. Students will first learn the foundations of deep learning, after which they will be introduced to a series of deep models: convolutional neural networks, autoencoders, recurrent neural network, and deep generative models. Students will work on case studies of deep learning in different fields such as computer vision, medical imaging, natural language processing, etc.
Goal	This course aims to familiarize students with the foundations of deep learning and its application domains. Specifically, students will be able to build, train, evaluate, and improve appropriate deep learning models for different problems.
Contents	This course covers the following major modules: (I) neural networks, (II) convolutional neural networks, (III) recurrent neural networks, (IV) reinforcement learning, and (V) generative adversarial networks.
Recommended Textbooks	I. Goodfellow, Y. Bengio, A. C. Courville, <i>Deep Learning</i> , MIT Press, 2016.
Recommended References & Supplemental Material	Relevant research papers, tech reports, and surveys for each topic, where needed, are identified in the teaching plan ahead. In addition, the following textbook may be useful: P. J. Gibson, A. Deep learning: A practitioner's approach, O'Reilly Media, Inc., 2017.



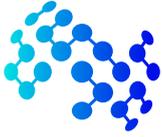
Teaching Week	Topics
1	Applied Math and Machine Learning Basics Lectures <ul style="list-style-type: none">• A refresher on linear algebra and probability• Machine learning basics Lab <ul style="list-style-type: none">• Getting started with Python• Practice problems to gain familiarity with Pytorch and Jupyter Notebook.
2	Introduction to Deep Learning and Neural Networks Basics Lecture <ul style="list-style-type: none">• History and cognitive basis of neural computation.• The neural network as a universal function approximator Lab <ul style="list-style-type: none">• Practice writing vectorized code with NumPy• Implementation and application of K-nearest neighbour algorithm on a toy dataset
3	Deep Feedforward Networks Lecture <ul style="list-style-type: none">• Artificial neural architecture design• Gradient-based learning• Optimization by gradient descent Lab <ul style="list-style-type: none">• Instructor-led demonstration related to topics of the week
4	Regularizing Neural Networks Lecture <ul style="list-style-type: none">• Regularization and under-constrained problems• Data augmentation• Early stopping• Dropout Lab <ul style="list-style-type: none">• Implement a deep fully connected neural network in the given python code template• Train it on a given toy dataset• Implement Dropout in the given python code template to regularize the deep fully connected network



Teaching Week	Topics
5	Optimizers for Training Deep Neural Networks Lecture <ul style="list-style-type: none">• Momentum, Nesterov momentum• Second order methods• Algorithms bearing adaptive learning rates• Hyperparameter optimization Lab <ul style="list-style-type: none">• Implement AdaGrad, RMSprop and ADAM optimizers in the given code skeleton
6	Convolutional Neural Networks (CNNs) - Part I Lecture <ul style="list-style-type: none">• Overview of CNNs• CNN basics – layer, kernel, stride, pooling, deconvolution, etc.• Convolution and pooling as an infinitely strong prior• Structured outputs Lab <ul style="list-style-type: none">• Instructor-led demonstration related to topics of the week
7	Convolutional Neural Networks (CNNs) - Part II Lecture <ul style="list-style-type: none">• Efficient convolution algorithms• Unsupervised features• Popular Image Classification CNNs Lab <ul style="list-style-type: none">• Implement the forward and backward passes of convolution, pooling, and batch normalization in the given code template
8	Recurrent Neural Networks (RNNs) - Part I Lecture <ul style="list-style-type: none">• Motivation for recurrent neural networks• Backpropagation through time (BPTT)• Architecture overview: fixed-size to sequence, sequence to sequence, etc. Lab <ul style="list-style-type: none">• Implement and train a vanilla RNN character-level language model



Teaching Week	Topics
9	Recurrent Neural Networks (RNNs) - Part II Lecture <ul style="list-style-type: none">• Long-term dependencies - challenges• Long short term memory units (LSTMs)• Gated recurrent units (GRUs)• Optimization for long-term dependencies Lab <ul style="list-style-type: none">• Implement and train LSTM character-level language model
10	Improving Deep Neural Networks Lecture <ul style="list-style-type: none">• Performance metrics• Default baseline models• Importance of data• Debugging strategies Lab <ul style="list-style-type: none">• Assignment # 3 due
11	Prevalent Deep Learning Fields Lecture <ul style="list-style-type: none">• Computer vision• Speech recognition, and natural language processing• Health care• Financial data analysis Lab <ul style="list-style-type: none">• Run any three state-of-the-art face detection algorithms and compare performance using standard metrics
12	Autoencoders Lecture <ul style="list-style-type: none">• Undercomplete autoencoders• Denoising autoencoders• Learning manifolds with autoencoders Lab <ul style="list-style-type: none">• Prepare a short report on the effectiveness of the U-Net architecture for both medical imaging applications and various mainstream computer vision tasks



Teaching Week	Topics
13	Generative Models Lecture <ul style="list-style-type: none">• Linear generative models• Generative adversarial networks (GANs) Lab <ul style="list-style-type: none">• Implement a GAN to generate photo-realistic (facial or natural) images
14	Deep Reinforcement Learning Lecture <ul style="list-style-type: none">• Introduction to reinforcement learning• Markov process• Value and policy iterations• Temporal difference (TD) learning• State-action-reward-state-action (SARSA) Lab <ul style="list-style-type: none">• Assignment # 4 due
15	Guest Lecture & Review Lecture <ul style="list-style-type: none">• Guest Lecture on Trending Topic Lab <ul style="list-style-type: none">• Review and Exam Preparation