

DS/CS 541 : Deep Learning

This course will offer a mathematical and practical perspective on artificial neural networks for machine learning. Students will learn about the most prominent network architectures including multilayer feedforward neural networks, convolutional neural networks (CNNs), auto-encoders, recurrent neural networks (RNNs), and generative-adversarial networks (GANs). This course will also teach students optimization and regularization techniques used to train them — such as back-propagation, stochastic gradient descent, dropout, pooling, and batch normalization. Connections to related machine learning techniques and algorithms, such as probabilistic graphical models, will be explored. In addition to understanding the mathematics behind deep learning, students will also engage in hands-on course projects. Students will have the opportunity to train neural networks for a wide range of applications, such as object detection, facial expression recognition, handwriting analysis, and natural language processing.

Department

Data Science

Computer Science

Credits 3.0

Prerequisites

Machine Learning (CS 539), and knowledge of Linear Algebra (such as MA 2071) and Algorithms (such as CS 2223)