

Neural Networks 1

Modality: Self-Paced Learning

Duration: 26 Hours

About this Course:

AI is revolutionizing the way we live, work and communicate. At the heart of AI is Deep Learning. Once a domain of researchers and PhDs only, Deep Learning has now gone mainstream thanks to its practical applications and availability in terms of consumable technology and affordable hardware.

The demand for Data Scientists and Deep Learning professionals is booming, far exceeding the supply of personnel skilled in this field. The industry is clearly embracing AI, embedding it within its fabric. The demand for Deep Learning skills by employers -- and the job salaries of Deep Learning practitioners -- are only bound to increase over time, as AI becomes more pervasive in society. Deep Learning is a future-proof career.

Within this series of courses, you'll be introduced to concepts and applications in Deep Learning, including various kinds of Neural Networks for supervised and unsupervised learning. You'll then delve deeper and apply Deep Learning by building models and algorithms using libraries like Keras, PyTorch, and Tensorflow. You'll also master Deep Learning at scale by leveraging GPU accelerated hardware for image and video processing, as well as object recognition in Computer Vision.

Throughout this program you will practice your Deep Learning skills through a series of hands-on labs, assignments, and projects inspired by real world problems and data sets from the industry. You'll also complete the program by preparing a Deep Learning capstone project that will showcase your applied skills to prospective employers.

This program is intended to prepare learners and equip them with skills required to become successful AI practitioners and start a career in applied Deep Learning.

Course Objectives:

- Fundamental concepts of Deep Learning, including various Neural Networks for supervised and unsupervised learning.
- Build, train, and deploy different types of Deep Architectures, including Convolutional Networks, Recurrent Networks, and Autoencoders.
- Application of Deep Learning to real-world scenarios such as object recognition and Computer Vision, image and video processing, text analytics, Natural Language Processing, recommender systems, and other types of classifiers.
- Master Deep Learning at scale with accelerated hardware and GPUs.
- Use of popular Deep Learning libraries such as Keras, PyTorch, and Tensorflow applied to industry problems.

Course Outline:

This Course Includes:

- Course Introduction
- Lecture 1: Neural Networks
- Lecture 2: Deep Neural Networks and Backpropagation
- Lecture 3: Modularity of Neural Networks
- Lecture 4: Gradient-based Learning
- Lecture 5: Demo: Building a Neural Network in Tensorflow