DR JUAN H KLOPPER
SCHOOL FOR DATA SCIENCE AND COMPUTATIONAL THINKING
STELLENBOSCH UNIVERSITY

The Fundamentals of Deep Learning

A course in the creation and use of deep neural networks in medicine and the life sciences using TensorFlow and Python

Introduction

This is a course that teaches the fundamentals of deep learning. Deep neural networks are algorithms that aim to mimic human learning using mathematical constructs.

Deep learning uses neural networks as models to predict outcomes or numerical values, having learnt from labelled data. As such, deep learning is a form of machine learning, which is in turn a form of artificial intelligence.

Machine learning has undergone a massive boon in the current age of big data and cheap computational power. Machines are now powerful enough to manage large amounts of data and the complicated mathematics required when using neural networks, random forests, nearest neighbours, support vector machines, and other learning tasks.

Python has become the leading computer language in the discipline of data science. TensorFlow, designed by Google, provides a framework in Python that makes deep learning a relatively simple task. The language is easy to learn and the TensorFlow framework does all the heavy mathematical lifting required by deep neural networks, behind the scenes.

This course starts by developing an intuitive understanding of the basics of the mathematics involved. It requires only high school mathematics, and the intuition is only built to provide for a better understanding of neural networks. The substance of the course is then about learning how to write short lines of Python code to construct neural networks. The course concentrates on two use cases.

Firstly, we look at building predictive models using structured data. This is the sort of data captured in a spreadsheet. We will use existing data from which our neural network will learn to predict an outcome variable. This is the same technology that predicts what you may want to watch next on your favourite streaming service or recommends what you might want to purchase on your favourite online shopping site. It can also learn to predict patient outcomes, having learnt from existing patient data.

The second case is that of computer vision. We will train a model to recognises malignant skin lesions from photos. This is a very common task in machine learning and deep neural networks are particularly adept at computer vision. They power self-driving cars after all.

Outcomes

By successfully completing this course you will have a deep appreciation of the inner workings of neural networks as they pertain to structured data and to images. You will be familiar with the different types of learning in artificial intelligence, how to work with data, how to create neural networks, how to train them on existing data, and how to test their accuracy.

Assessment

There will be no assessment of any mathematics in this course. Instead, your ability to write Python code to generate deep neural networks will be assessed. You will have to create and test your own neural networks. Fortunately, this is easily achieved using Python.

System requirements

There will be no need to install Python on your system. For those who are curious to do so, information will be provided. The course will be run using Google Colaboratory. You will have to have access to the internet and will have to create a Google account if you do not already have one.

Google Colaboratory looks very much like Google Docs. It is a blank canvas on which we write normal sentences just as in Google Docs. It also allows us to write Python code. Google provides the hardware required to store the data that we will use and the horsepower to train our neural networks, all without any cost.

Course syllabus

This course will cover the following 15 sections:

1. The basics of Google Colaboratory and Python
    a. In this section you will learn how to open a Colaboratory (Jupyter) notebook and how to use it to generate documents that contain normal English text and Python code
    b. You will also learn the basics of the Python Language
2. An introduction to deep learning
    a. This section provides the following…
        i. An overview of the history of artificial intelligence
        ii. An intuitive understanding of linear regression to serve as the basic building block of a learning model

        iii. The concept of a loss function and how it is the quintessential construct in the learning process
        iv. How data flows through a neural network
        v. How the neurons in a neural network update their state to improve their predictive power
3. Introductory mathematics
    a. This section covers the fundamentals of linear algebra and multivariable differentiation
    b. It is an optional section and will not be assessed
    c. It does provide for clear intuition of the inner working of the learning process in neural networks
4. Forward propagation and backpropagation
    a. This section provides a deeper look at how data flows through a neural network and how the nodes (neurons) are updated to improve their learning ability
    b. It is an optional section and will not be assessed
5. Generating simulated data and craeting a neural network
    a. In this section we look at how to generate simulated data, ideal for use when learning the basics of neural networks
    b. We also create our first neural network
6. A deeper dive into a network that uses structured data
    a. In this section we create a more robust neural network
7. Bias and variance
    a. In this section we learn about the accuracy of a deep neural network and how to balance the learning process so as to perform well on unseen data
8. Improving performance
    a. In this section we improve the performance of our neural network by employing a variety of learning techniques
9. Convolutional neural networks
    a. In this section we learn how to construct neural networks that can recognize images
10. Learning to classify malignant skin lesions
    a. In this capstone section, we create a neural network that can distinguish between benign and malignant skin lesions

Duration and dates

This course will be presented over five days.

- Day 1
    - Sections 1 and 2
- Day 2
    - Sections 3 and 4
- Day 3
    - Sections 5 and 6
- Day 4
    - Sections 7 and 8

- Day 5
    - Sections 9 and 10

Sessions

The five days will be divided as below. Adequate time will be provided for breaks. The course material will be available beforehand, and it is advised that preparation for the following day be done every evening.

Formal lectures will contain a live session and recorded video.

All assessments must be forwarded by 14H00 every day.

09h00-12h00
Formal lecture

12h00-14h00
Complete assessments

14h00-16h00
Discussion of assessments

Students are required to attend all lectures and submit assessments in order to receive the attendance certificate.