
DEEP LEARNING ONLINE BOOTCAMP

DETAILED SYLLABUS: BEGINNERS

August 20 - September 06

Overview

In our endeavour to build data culture and democratize Data Science learning, we are launching a deep learning bootcamp with the help of academia and industry experts. The online bootcamp will have a series of day-wise learning modules/challenges and live sessions by data science experts.

In this bootcamp, you'll be able to grasp the fundamentals of deep learning that will aid your Data Science journey and get you started with building the required industry skills.

This is a community initiative, driven by experts and mentors, and you have the opportunity to attend it for free.

Prerequisites

- Python
- Numpy
- Pandas
- Linear Algebra
- Jupyter Notebook/ Google Colab

Module 0: Data Science Fundamentals

Refresh your data science basics that will be useful across the whole Bootcamp.

LEARNING OUTCOMES

Day 0 Aug. 19	Data Science Refresher	<ul style="list-style-type: none">• Numpy and Pandas fundamentals• Linear Algebra fundamentals• Jupyter notebook/Anaconda/Google Colab• Visualization fundamentals
------------------	-------------------------------	---

Module 1: Deep Learning and Neural Networks Basics

Dive into the world of deep learning and understand what a Neural Network is

LEARNING OUTCOMES

Day 1	Overview	<ul style="list-style-type: none"> • What is DL and why is it so popular? • Comparison between ML and DL • Real-life applications of DL • Deep Learning Frameworks • Refresher of Machine Learning Concepts
Day 2	Introduction to the Simplest Neural Network: Perceptron	<ul style="list-style-type: none"> • Neural Networks: The heart of DL • What is NN? • Neurons: The building block of NN • Single and Multilayer Perceptrons (MLP) • Similarity to the brain neurons
Day 3	Introduction to tf.Keras	<ul style="list-style-type: none"> • What is Keras? • Keras vs tf.Keras • How is Tensorflow relevant in the industry? • How to Install TensorFlow? • How to Confirm TensorFlow Is Installed? • Tensors • Playing around with some Tensor operations • Basic Introduction to GPUs and CUDA

Module 2: Neural networks for regression problems + Hyperparameter Tuning

Get an overview of how a Regression model can be built with tf.Keras and learn to perform Hyperparameter Tuning

LEARNING OUTCOMES

Day 4	Overview	<ul style="list-style-type: none">• Regression model with tf.Keras• Epochs• Learning Rate• Batch Size• Hyperparameter Tuning
-------	-----------------	--

Module 3: Neural networks for classification problems

Build Classification Models on diverse datasets

LEARNING OUTCOMES

Day 5	Multi-Layer Perceptron/ Deep NN	<ul style="list-style-type: none"> ● Neural Network Architecture ● NN Working: <ul style="list-style-type: none"> ○ Feedforward intuition ○ Gradient Descent ○ Backpropagation intuition ● The 5 step model life cycle
Day 6	MLP for Binary Classification	<ul style="list-style-type: none"> ● Building a Deep Learning Model: <ul style="list-style-type: none"> ○ The basic difference between Sequential and Functional Model API ○ Building an MLP for binary classification using Sequential Model ● Compiling a model <ul style="list-style-type: none"> ○ Loss Functions ○ Optimizers ● Activation functions
Day 7	MLP for Multiclass Classification	<ul style="list-style-type: none"> ● Building a Deep NN on MNIST Dataset <ul style="list-style-type: none"> ○ Pre-processing techniques ○ Deciding Loss and Optimization Functions ● Train, Validation and Test Set ● Training the model ● Hyperparameter Tuning ● Testing the model ● Saving and Loading Models

Module 4: Optimizing a Neural Network

Optimize the training of your Neural Network to create refined and accurate models

LEARNING OUTCOMES

Day 8	Optimising the training of a NN - part 1	<ul style="list-style-type: none">• Overfitting and Underfitting• Early Stopping• Regularization
Day 9	Optimising the training of a NN - part 1	<ul style="list-style-type: none">• Dropout• Local and Global Minima• Introducing the final assignment
Day 10	Optimising the training of a NN - part 2	<ul style="list-style-type: none">• Batch vs Stochastic Gradient Descent• Random Initialization/Restart• Vanishing Gradient• Other Activation Functions

Module 5: **Applied Problem Solving**

Datathon: Get your hands dirty with applied problem solving