Deep learning with Tensorflow

Overview:

TensorFlow is one of the most in-demand and popular open-source deep learning frameworks available today.

The DeepLearning.AI TensorFlow Developer Professional Certificate program teaches you applied machine learning skills with TensorFlow so you can build and train powerful models. In the DeepLearning.AI TensorFlow Developer Professional Certificate program.

By the end of this program, you will be ready to:

- Build and train neural networks using TensorFlow
- Improve your network’s performance using convolutions as you train it to identify real-world images
- Teach machines to understand, analyse, and respond to human speech with natural language processing systems
- Process text, represent sentences as vectors.

Content:

Introduction

- Why you should study machine learning?
- What does the course cover?

Introduction to neural networks

- Introduction to neural networks
- Training the model
- Types of machine learning
- The linear model
- Need Help with Linear Algebra?
- The linear model. Multiple inputs
- Graphical representation
- The objective function
- L2-norm loss
- Cross-entropy loss
- One parameter gradient descent
- N-parameter gradient descent

Setting up the working environment

- Setting up the environment - An introduction
- Installing Anaconda
- Jupyter Shortcuts
- Installing TensorFlow 2
- Installing packages - exercise
- Installing packages - solution
- Minimal example - Exercises

TensorFlow - An introduction

- TensorFlow outline
- TensorFlow intro
- A Note on Coding in TensorFlow
- Types of file formats in TensorFlow and data handling
- Model layout - inputs, outputs, targets, weights, biases, optimizer and loss
- Interpreting the result and extracting the weights and bias
- Customizing your model
- Minimal example with TensorFlow - Exercises

Going deeper: Introduction to deep neural networks

- Layers
- What is a deep net?
- Understanding deep nets in depth
- Why do we need non-linearities?
- Activation functions
- Softmax activation
- Backpropagation
- Backpropagation - visual representation
- Backpropagation. A peek into the Mathematics of Optimization

Overfitting

- Underfitting and overfitting
- Training and validation
- Training, validation, and test
- N-fold cross validation
- Early stopping

Initialization

- Initialization - Introduction
- Types of simple initializations
- Xavier initialization

Gradient descent and learning rates

- Stochastic gradient descent
- Gradient descent pitfalls
- Momentum
- Learning rate schedules
- Adaptive learning rate schedules
- Adaptive moment estimation

Pre-processing

- Pre-processing introduction
- Basic pre-processing
- Standardization
- Dealing with categorical data
- One-hot and binary encoding

The MNIST

- The dataset
- How to tackle the MNIST
- Importing the relevant packages and load the data
- Pre-process the data - create a validation dataset and scale the data
- Pre-process the data - scale the test data
- Pre-process the data - shuffle and batch the data
- Pre-process the data - shuffle and batch the data
- Outline the model
- Select the loss and the optimizer
- MNIST
- Testing the model

Business case

- Exploring the dataset and identifying predictors
- Outlining the business case solution
- Balancing the dataset
- Pre-processing the data
- Pre-processing exercise
- Load the pre-processed data
- Load the pre-processed data - Exercise
- Learning and interpreting the result
- Setting an early stopping mechanism
- Setting an early stopping mechanism - Exercise
- Testing the model

Appendix: Linear Algebra Fundamentals

- What is a Matrix?
- Scalars and Vectors
- Linear Algebra and Geometry
- Scalars, Vectors and Matrices in Python
- Tensors
- Addition and Subtraction of Matrices
- Errors when Adding Matrices
- Transpose of a Matrix
- Dot Product of Vectors
- Dot Product of Matrices
- Why is Linear Algebra Useful?

Conclusion

- What's further out there in the machine and deep learning world
- An overview of CNNs
- How DeepMind uses deep learning
- An overview of RNNs
- An overview of non-NN approaches