Course Outline

AI for Image Processing

Course Summary

Description

In the last few years, AI algorithms for image analysis have made tremendous progress. This is mainly due to the abundance of data, affordable computing, and exceptional libraries. Google has open-sourced a library called TensorFlow which has become the de facto standard, allowing the state of the art machine learning done at scale, complete with GPU based acceleration.

This course introduces Deep Learning concepts and TensorFlow and Keras libraries to students.

Objectives

At the end of the course, students will understand:

- Deep Learning concepts.
- TensorFlow and Keras.
- Create ceural networks with Tensorflow and Keras.
- Learn to use tools like Tensorboard to help with training neural networks.
- Deep Neural Networks.
- Convolutional Neural Networks (CNN).
- Generative Adversarial Networks (GAN).
- Auto Encoders.

Topics

- Introduction to Deep Learning
- Introducing TensorFlow
- Introducing Keras
- Deep Learning Concepts
- Feedforward Network
- Computer Vision
- Generative Adversarial Network (GAN)
- Auto Encoder
- Transfer Learning

Audience

The audience for this class includes Developers, Data Analysts, and Data Scientists.

Prerequisites

Attendees should possess basic knowledge of Python and Jupyter notebooks.

Duration

Three days

AI for Image Processing

Course Outline

I. Introduction to Deep Learning
- A. Understanding Deep Learning use cases
- B. Understanding AI/Machine Learning/Deep Learning
- C. Data and AI
- D. AI vocabulary
- E. Hardware and Software Ecosystem
- F. Understanding types of Machine Learning (Supervised/ Unsupervised/Reinforcement)

II. Introducing Tensor Flow
- A. Tensorflow intro
- B. TensorFlow features
- C. Execution graph
- D. TensorFlow on GPU and TPU
- E. TensorFlow API
- F. Lab: Setting up and Running TensorFlow

III. Introducing Keras
- A. Keras Intro
- B. Keras conceps (models, layers)
- C. Using Keras API
- D. Lab

IV. Deep Learning Concepts
- A. Introducing Perceptrons
- B. Linear Perceptrons
- C. Activation Functions (Sigmoid, Tanh, Relu, Softmax)
- D. Backpropagation
- E. Optimizers (Gradient Descent, Adam, RMSProp)
- F. Loss functions for regressions and classifications
- G. Vanishing/exploding gradient problem
- H. Lab: TensorFlow playground

V. Feedforward Network
- A. FFNN architecture
- B. Input layer, output layer
- C. Hidden layers and Deep neural networks
- D. Sizing neural networks
- E. Lab: Feedforward Neural Networks

VI. Computer Vision
- A. Introducing Convolutional Neural Networks (CNN)
- B. CNN architecture
- C. CNN concepts
- D. Lab: Image recognition using CNNs

VII. Generative Adversarial Network (GAN)
- A. GAN Overview
- B. Generating Images
- C. Lab: GAN lab

VIII. Auto Encoder
- A. Auto Encoder overview
- B. Auto Encoder use cases
- C. Lab: Auto encoder

IX. Transfer Learning
- A. Understanding transfer learning
- B. Customizing available models
- C. Lab: transfer learning lab
- D. Lab: Benchmark performances on CPU and GPU