AI Courses by OpenCV

DEEP LEARNING WITH TENSORFLOW & KERAS

Getting Started

1. Introduction to Artificial Intelligence

    - History of AI
    - Applications of AI
    - AI in Computer Vision
    - AI Terminology
    - Introduction to Deep Learning
    - Deep Learning Frameworks

2. NumPy Refresher

    - NumPy Refresher Part-1
    - NumPy Refresher Part-2
    - NumPy Refresher Part-3

3. Introduction TensorFlow
4. What is inside an ML Algorithm

    - Machine Learning pipeline
    - Solving ML Problems
    - Gradient Descent for Optimization

5. Regression: A Classic Supervised Learning Problem

Assignment 1: Implement Leaky ReLU, Softmax and Convolution using TensorFlow

Assignment 2: Implement Gradient Descent for two variables

Module 2 : Neural Networks

1. Understanding Neural Networks

    - Deep Learning Overview
    - What is a Neural Network
    - Feature Vectors and Normalization
    - Demystifying Neural Networks

2. Building Neural Network in Keras

    - Introduction to Linear Regression
    - Auto-MGP Data Processing
    - Linear Regression with Keras
    - Binary Classification with Keras

3. Building Blocks of a Neural Network

    - Loss Function for Regression
    - Loss Function for Classification
    - Types of Activation Functions
    - How does the network learn?

4. Multi-class Classification using Keras

    - Classifying MNIST digits with a Multi-Layer Perceptron (MLP)

5. Model Complexity, Generalization and Handling Overfitting

    - Bias Variance Trade-off
    - How to Prevent Overfitting

Assignment 3: Applying a MLP on the Fashion MNSIT Dataset

Module 3 : Convolutional Neural Network

1. Image Classification

    - Image classification using CNN

2. CNN

    - CNN Building Blocks
    - The Convolution Operation
    - Layers in CNN
    - Implementing LeNet in Keras

3. Evaluation metrics for Classification

    - Performance Metrics for Classification
    - Evaluation metrics for Classification

4. Building Models with Custom Data

    - Keras Image _Dataset_from_Directory
    - Overfitting and Data Augmentation

5. Working with pretrained Networks

    - Important CNN Architectures
    - Pretrained Models for Keras Applications
    - Training VGGNet from Scratch on Balls Dataset

6. Transfer Learning and Fine-Tuning

    - Transfer Learning with VGGNet as Feature Extractor on Balls Data
    - Transfer Learning with VGGNet as Feature Extractor on ASL Data
    - Fine Tuning VGGNet using ASL Data

Assignment 4: Sequential vs Functional API

Assignment 5: Image Classification using CNN

Project1: Implement an Image Classifier from scratch

Module 4 : Advance Training Concepts

1. Optimizers
2. Handling Data in TensorFlow
    - Introduction to TF Data
    - Custom Data Loader using Sequence Model
    - TF Records
3. Learning Rate Schedulers
    - Learning Rate Decay Models
    - LR Schedulers
4. Gaining Insights
    - GradCam

Project 3: Kaggle Competition - Classification

Module 5 : Semantic Segmentation

5. Introduction to Semantic Segmentation
    - Introduction to Semantic Segmentation
    - Semantic Segmentation Datasets
    - Overview of Semantic Segmentation
6. Custom Data Loader
    - Introduction to Segmentation Datasets and Custom Data Loader
7. Transposed Convolution
8. Fully Convoluted Networks

- FCN Architecture
- FCN on Raod Data: CE Loss
- Evaluation Metrics in Semantic Segmentation
- FCN: Custom Metrics and Loss Functions

9. Evaluation Metrics for Semantic Segmentation
10. UNet

- UNet Architecture
- UNet on Road Data: CE Loss
- UNet on CamVid Data: Dice Loss

11. Dilated Convolutions
12. DeepLabV3

- DeepLabv3 Architecture
- DeepLabv3+ on Road Data: CE Loss
- DeepLabv3+ on CamVid Data: Dice Loss
- DeepLabv3+ on CamVid Data: Best Results
- DeepLabv3+ on SUIM Data: CE Loss
- DeepLabv3+ on SUIM Data: Best Results

Assignment 6: PSPNet

Project 3: Semantic Segmentation

Module 6 : Object Detection

1. Introduction to Object Detection

   - History of Object Detection
   - Object Detection Datasets

2. Hands on with Object Detector

   - Inference using Object Detection Models from TensorflowHub

3. Classification to Detection

    - Image Classification vs Object Detection
    - Revisiting Classification Pipeline
    - Encoding Bounding Boxes using Anchors
    - IoU
    - Encoding of Ground Truth
    - Multiple Anchors

4. Non Maximum Suppresion (NMS)

    - Introduction to NMS
    - NMS vs Soft NMS

5. Evaluation Metrics

    - Why we need Evaluation Metrics
    - Building Blocks of mAP
    - Precision and Recall
    - Average Precision (AP) and Mean Average Precision (mAP)

6. Popural Object Detection Architectures

    - Traditional Object Detectors
    - Two Stage Object Detectors
    - YOLO: You Only Look Once
    - SSD: Single Shot MultiBox Detector
    - RetinaNet

7. TensorFlow Object Detetcion API

    - Installation of TFOD
    - Introduction and Infernece using TFOD Pretrained Models
    - Data Preparation in TFOD
    - Pipeloine Configuration in TFOD
    - Inference with a Pretrained Model

8. Create a Custom Object Detetcor

- Detetcor Architecture
- Anchor Boxes and Label Encoding
- Anchors Generatio using Keras
- Loss Function
- Decode NMS
- Evaluator in the Pipeline
- Create a Custome Data Loader
- Trainign from Scratch

Assignment 7: Encoding and Decoding of Ground Truths for Anchor box implementation

Project 4: Object Detection

Module 7 : Introduction to Generative Adversarial Networks

1. Introduction to GANs
2. Vanilla GAN using Fashion MNIST
3. DCGAN using Flickr Faces
4. CGAN using Fashion MNIST

Module 8 : Introduction to Mediapipe and Applications

1. Introduction to Mediapipe
2. Posture Analysis using Mediapipe
3. Drowsy Driver Detection using Mediapipe