

[2021 Update] Deep Learning Specialization Change Log

Overview

This document contains a detailed changelog for the Deep Learning Specialization from DeepLearning.AI. This Specialization was updated in April 2021 to satisfy the new deployment options available on Coursera and to include developments in deep learning and programming frameworks that have occurred after the initial release in 2017.

All changes are described in detail and color-coded:

- Orange: Changes to the autograders
- Blue: Updated from Tensorflow 1.13 to Tensorflow 2.3
- Green: Completely new course material

[Course 1] Neural Networks and Deep Learning

- **Week 1:** Introduction to Deep Learning
- **Week 2:** Neural Network Basics
 - Notebook: Python Basics with NumPy
 - **New autograder: nbgrader**
 - Notebook: Logistic Regression with a Neural Network mindset
 - **New autograder: nbgrader**
- **Week 3:** Shallow Neural Networks.
 - Notebook: Planar data classification with a hidden layer
 - **New autograder: nbgrader**
- **Week 4:** Deep Neural Networks.
 - Notebook: Building your deep neural network: Step by Step
 - **New autograder: nbgrader**
 - Notebook: Deep Neural Network - Application
 - **New autograder: nbgrader**

[Course 2] Improving Deep Neural Networks: Hyperparameter tuning, Regularization, and Optimization

- **Week 1:** Practical Aspects of Deep Learning
 - Notebook: Initialization
 - New autograder: nbgrader
 - Notebook: Regularization
 - New autograder: nbgrader
 - Notebook: Gradient Checking
 - New autograder: nbgrader
- **Week 2:** Optimization Algorithms
 - Notebook: Optimization Methods
 - New autograder: nbgrader
 - Add learning rate scheduling and decay
- **Week 3:** Hyperparameter Tuning, Batch Normalization, and Programming Frameworks
 - New Video: Tensorflow 2.3 demo
 - Notebook: TensorFlow tutorial
 - New autograder: nbgrader
 - A brief introduction to tf.datasets and GradientTape

[Course 3] Structuring Machine Learning Projects

- **Week 1:** ML Strategy (1)*
 - **Week 2:** ML Strategy (2)*
- *Note: No changes as there are no coding assignments in this course*

[Course 4] Convolutional Neural Networks

- **Week 1:** Foundations of Convolutional Neural Networks
 - Notebook: Convolutional Model: step by step
 - Notebook: Convolutional Model: Application
 - Merged this assignment with the Keras intro from W2. Here the student will learn how to use the Sequential API (using HappyHouse dataset) and the Functional API (using SIGNS dataset) of TensorFlow.
 - New autograder: nbgrader
- **Week 2:** Deep Convolutional Models: case studies
 - Notebook: Residual Networks

- New autograder: nbgrader
 - New notebook on **transfer learning**
 - Uses **MobileNet** as a lightweight model to fine-tune.
 - New autograder: nbgrader
 - New videos: MobileNet and EfficientNet
 - MobileNet
 - MobileNet Architecture
 - EfficientNet
- **Week 3: Object Detection**
 - Notebook: Car detection with YOLO
 - New autograder: nbgrader
 - New notebook: Application to semantic segmentation.
 - Uses **U-Net** to make image segmentation in a self-driving cars dataset (CARLA)
 - New autograder: nbgrader
 - New video(s) for Semantic Segmentation with U-Net
 - Semantic Segmentation with U-Net
 - Transpose Convolutions
 - U-Net Architecture Intuition
 - U-Net Architecture
- **Week 4: Special applications: Face recognition & Neural style transfer.**
 - Notebook: Neural Style Transfer:
 - Uses GradientTape to keep track of gradients.
 - New autograder: nbgrader
 - Notebook: Face recognition
 - New autograder: nbgrader

[Course 5] Sequence Models

- **Week 1: Recurrent Neural Networks**
 - Notebook: Building a recurrent neural network - step by step.
 - New autograder: custom grader
 - Notebook: Dinosaur Island - Character-Level Language Modeling
 - New autograder: custom grader
 - Notebook: Jazz improvisation with LSTM
 - New autograder: custom grader
- **Week 2: Natural Language Processing & Word Embeddings**

- Notebook: Operations on word vectors - Debiasing
 - New autograder: custom grader
 - Notebook: Emojify
 - New autograder: custom grader
- **Week 3: Sequence models & Attention mechanism**
 - Notebook: Neural Machine Translation with Attention
 - New autograder: custom grader
 - Notebook: Trigger word detection
 - New autograder: custom grader
- **Week 4: Transformer Network**
 - New Notebook: Self-Attention & Transformers
 - Build the basic transformer architecture
 - New autograder: custom grader
 - New ungraded lab: Transformer applications: Named Entity Recognition
 - Fine-tune pre-trained transformer models to your data to solve the NER problem
 - New ungraded lab: Transformer applications: Question Answering
 - Fine-tune pre-trained transformer models to your data to solve the QA problem
 - New videos for transformer architecture
 - Transformer Network Intuition
 - Self-Attention
 - Multi-Head Attention
 - Transformer Network