ACC: ACC 3

ACC Bucket : Python & Machine Learning - 3

Course: Applied Deep Learning

OVERALL COURSE OBJECTIVES: The objective of this course is to synthesize comprehensive deep learning knowledge, cultivate capabilities to configure and optimize neural networks, refine machine learning troubleshooting strategies, and apply advanced concepts, all while actively contributing to the development of leading-edge AI technology.

LEARNING OUTCOMES: On successful completion of the course the students shall be able to:

1. Construct, train, and implement fully connected deep neural networks and apply them to various applications.
2. Evaluate and employ standard neural network techniques such as initialization, L2 and dropout regularization, hyperparameter tuning, batch normalization, and gradient checking.
3. Implement a variety of optimization algorithms, including mini-batch gradient descent, Momentum, RMSprop and Adam, and check for their convergence.
4. Implement a neural network in TensorFlow.
5. Diagnose errors in a machine learning system and strategize solutions for reducing these errors.
6. Apply advanced machine learning concepts such as end-to-end learning, transfer learning, and multi-task learning.

COURSE CONTENT:

Module 1: Neural Networks and Deep Learning [25 Hours]

In this foundational course of the Deep Learning Specialization, you will delve into neural networks and deep learning. The course aims to familiarize you with significant technological trends propelling deep learning, enabling you to build, train, and apply fully connected deep neural networks. It will also guide you to implement efficient (vectorized) neural networks, identify key parameters of a neural network's architecture, and apply deep learning to your applications. This specialization forms a cornerstone for understanding the potential, challenges, and implications of deep learning. It creates an avenue for you to gain the proficiency to employ machine learning in your work, elevate your technical career, and make significant strides in the AI landscape.

Sub-Topic
Supervised Learning with Neural Networks
Binary Classification
Derivatives with a Computation Graph
Vectorizing Logistic Regression
Shallow Neural Networks
Building Blocks of Deep Neural Networks

Formative Assessments:
4 quizzes and 4 graded coding/lab assignments.

Module 2: Improving Deep Neural Networks: Hyperparameter Tuning, Regularization and Optimization [24 Hours]

The second course in the Deep Learning Specialization delves into understanding the mechanics behind deep learning performance and consistently generating good results. Upon completion, you will master best practices for training, developing test sets, and analyzing bias/variance for deep learning applications. Additionally, you will acquire skills in implementing standard neural network techniques such as initialization, L2 and dropout regularization, hyperparameter tuning, batch normalization, gradient checking, and a variety of optimization algorithms. The course also covers implementing a neural network in TensorFlow. This specialization is foundational and prepares you for participation in advanced AI technology development, enabling you to apply machine learning practically, accelerate your technical career, and make significant strides in the AI domain.

Sub-Topic
Practical Aspects of Deep Learning
Understanding Exponentially Weighted Averages
Weight Initialization for Deep Networks
The Problem of Local Optima
Optimization Algorithms
Hyperparameter Tuning, Batch Normalization and Programming Frameworks

Formative Assessments:
3 quizzes and 5 coding/lab assignments

Module 3: Structuring Machine Learning Projects [7 Hours]

The third course of the Deep Learning Specialization aims to equip you with the ability to successfully execute a machine-learning project, cultivating your decision-making skills as a project leader. You will acquire the ability to diagnose errors in a machine learning system, prioritize error-reducing strategies, comprehend complex ML settings, and apply end-to-end learning, transfer learning, and multi-task learning. This course is suitable for learners with basic machine learning knowledge and offers industry experience drawn from Andrew Ng's extensive deep learning product development. It's perfect for those aspiring to become technical leaders in AI teams. The Deep Learning Specialization as a whole provides a sturdy foundation on which you can explore the capabilities, challenges, consequences of deep learning, and apply machine learning to your work, thus enhancing your technical career and consolidating your steps in the AI world.

Sub-Topic
Why ML Strategy
Understanding Human-level Performance
Surpassing Human-level Performance
Transfer Learning
Multi-task Learning
What is End-to-end Deep Learning?

Formative Assessments:

2 quizzes and 5 coding/lab assignments

ASSESSMENT:
For summative assessments, Coursera will provide question banks for which exams can be conducted on the Coursera platform or the faculty will create their own assessments.

Note: If a Course or Specialization becomes unavailable prior to the end of the Term, Coursera may replace such Course or Specialization with a reasonable alternative Course or Specialization.