Caltech | Center for Technology & Management Education

California Institute of Technology

Deep Learning Using TensorFlow

2022–2023

Syllabus

==========

Course Page: https://ctme.caltech.edu/deep-learning-open

What is Deep Learning?

Deep Learning is a branch of Artificial Intelligence (AI) based on the architecture of Neural Networks. When the number of hidden layers in a neural network is extended, it becomes a 'Deep Learning' Neural Network.

Computer algorithms are developed by specifying the rules of the problems and defining step-by-step instructions on how to solve a problem. When the problem is complex enough that the rules cannot be defined, an enormous amount of data, along with the answers, are fed into neural networks. Deep Neural Networks develop the algorithm themselves and then develop the ability to provide an answer for a new element.

Applications of Deep Learning

The applications of Deep Learning include object recognition in images, natural language processing, and human speech recognition. From a product development perspective, autonomous cars are likely to be a major application of Deep Learning. Object recognition within an image is a necessary aspect of such an application.

Deep Learning applications are found in virtually all industries (manufacturing, pharmaceutical, medical, security, government etc.). It is expected that there will be an extremely high demand for people with Artificial Intelligence, Neural Networks and Deep Learning expertise.

Tools

TensorFlow/Keras: Google has released many tools for the Deep Learning market. The most prominent one is TensorFlow, which was released in 2015. TensorFlow is Python-based low-level library, which is not very user friendly. Google has released another package called Keras, which is much more user friendly and provides user interface to the TensorFlow software.

Colab/Google Cloud Platform (GCP): Google's Colab provide Python/Jupyter environment to write Deep Learning Models. Colab communicates with Google Cloud Platform using the APIs for Artificial Intelligence services to build Deep Learning Models. TensorFlow/Keras packages can be accessed from Colab.

GPU/TPU: Google and Nvidia have introduced processors for the Artificial Intelligence (AI) market, which are called GPU (Graphics Processing Units) and TPU (Tensor Processing Units). These processors are extremely fast and can execute trillions of instructions per second in parallel. The architecture of the processors is well-suited for solving the problems for Deep Learning. Future Deep Learning applications are expected to use these processors with TensorFlow software to solve overly complex real-world problems.

Course Description

Course Description

Day #1

This course will first cover the concepts of Neural Networks and Deep Learning. It will cover the basic of the TensorFlow and Keras software packages. Google's Colab will be used as a platform for programming in Python/Jupyter. Neural Network mathematics will be covered in detail.

Day#2

Machine learning models for estimation (regression) and classification will be built using TensorFlow and Keras.

Day#3

Gradient Descent (GD) optimization is a technique that has been used successfully in many Machine Learning models. Details of Gradient Descent and Back Propagation algorithms will be covered. However, the current GD algorithm is slow to converge in Deep Learning models. Latest Deep Learning optimization techniques will be covered which includes Momentum, AdaGrad, RMS Prop and Adam. These optimization techniques overcome the shortcomings of the GD algorithm.

Day#4

Generative Neural Networks generate new images, text, and audio. These networks are the foundation of Speech and Translation services offered by Artificial Intelligence features of Cloud services. A Generative Adversarial Network (GAN) consists of two neural networks contest with each other in a game. GAN trained on photographs can generate new photographs that look at least superficially authentic to human observers, having many realistic characteristics. GAN technology is used for generating fake images and videos. This course will cover Generative Neural Networks.

Day#5

Machine Learning and Deep Learning models are currently built using TensorFlow/Keras. To train a Deep learning model requires a significant amount of computing resources. This precludes running Deep Learning models on personal computers. Generally, they are not powerful enough. Convolutional Neural Networks (CNN) and Generative Adversarial Networks (GAN), a specific type of Neural Networks, can only run on cloud servers because they have enough computing resources. The 3 most popular cloud services are AWS (Amazon Web Services), GCP (Google Cloud Platform), and Azure from Microsoft.

The GCP platform will be used to build the "Computer Vision" applications which include Object Detection, Image Classification, and Video Intelligence.

Day#6

The GCP platform will be used to build the Translation, Speech-to-Text, BigQuery and Natural Language applications.

Text Books

1 Optional	Neural Networks & Deep Learning by Charu Aggarwal Publisher: Springer	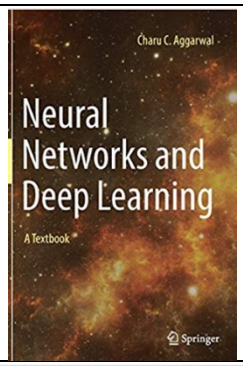
2 Optional	"Deep Learning" by Ian Goodfellow, Yoshua Bengio, and Aaron Courville	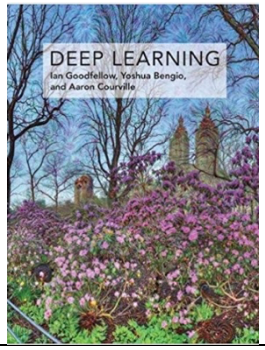

Lesson Plan

	Lesson Content
1	• Machine Learning, Neural Networks and Deep Learning • TensorFlow/Keras Architecture • Neural Network and NN Math
2	• Linear Regression in TensorFlow • Classification in TensorFlow
3	• Gradient Descent + Backpropagation Algorithms • Optimization: Momentum + Adaptive Gradient

	Lesson Content
4	• Convolutional Neural Networks • Generative Neural Networks + Autoencoders • Generative Adversarial Networks (GAN)
5	• Google Cloud Platform (GCP) Account Setup • GCP: Vision API: Object Detection + Image Classification • GCP: Video Intelligence
6	• GCP: Translation API • GCP: Converting Speech-to-Text API • GCP: BigQuery • GCP: Natural Language API

Course Objectives

At the end of this course, participants will be able to:

- Understand the Neural Networks and Deep Neural Networks
- Understand TensorFlow and Keras Architecture
- Build neural networks models using TensorFlow/Keras in Colab
- Implement Linear Regression, and Classification models in TensorFlow

- Implement Gradient Descent algorithm for optimization
- Implement Backpropagation algorithm in TensorFlow
- Understand Momentum, AdaGrad, RMS Prop and Adam Optimization algorithms
- Understand Autoencoders and Generative Neural Networks

- Setup Google Cloud Platform (GCP) account
- Implement CNN in TensorFlow and GCP using Vision API
- Implement Speech conversion into text application using GCP Speech API
- Implement Translate application to translate from one language to another

Instructor Information

Dr. Ash Pahwa
Office Phone: (949) 378-1229
Email: ash@ashpahwa.com
Website: www.AshPahwa.com

Ash Pahwa, Ph.D., is an educator, author, entrepreneur, and technology visionary with three decades of industry and academic experience. He has founded several successful technology companies during his career, the latest of which is A+ Web Services.

Dr. Pahwa earned his doctorate in Computer Science from the Illinois Institute of Technology in Chicago. He is listed in *Who's Who in the Frontiers of Science and Technology*. He is also a Google Certified Analytics Consultant. His expertise includes search engine optimization, web analytics, web programming, digital image processing, database management, digital video, and data storage technologies.

In Industry, Dr. Pahwa has worked for General Electric, AT&T Bell Laboratories, Xerox Corporation, and Oracle. He founded CD-Gen, Inc. and DV Studio Technologies, LLC., which introduced successful products for CD-Recording (CDR) and MPEG encoding. His book, *CD-Recordable Bible* was published in English, Japanese, and German.

In Academia, Dr. Pahwa teaches courses at California Institute of Technology (Pasadena) and the University of California system. Since 2008, he taught many courses at UC Irvine, UCLA, and UC San Diego.