Nvidia - Fundamentals of Deep Learning

Course 11214 – 8 Hours

Overview

Businesses worldwide are using artificial intelligence (AI) to solve their greatest challenges. Healthcare professionals use AI to enable more accurate, faster diagnoses in patients. Retail businesses use it to offer personalized customer shopping experiences. Automakers use it to make personal vehicles, shared mobility, and delivery services safer and more efficient. Deep learning is a powerful AI approach that uses multi-layered artificial neural networks to deliver state-of-the-art accuracy in tasks such as object detection, speech recognition, and language translation. Using deep learning, computers can learn and recognize patterns from data that are considered too complex or subtle for expert-written software.
In this workshop, you'll learn how deep learning works through hands-on exercises in computer vision and natural language processing. You'll train deep learning models from scratch, learning tools and tricks to achieve highly accurate results. You'll also learn to leverage freely available, state-of-the-art pre-trained models to save time and get your deep learning application up and running quickly.

By participating in this workshop, you'll:

- Learn the fundamental techniques and tools required to train a deep learning model
- Gain experience with common deep learning data types and model architectures
- Enhance datasets through data augmentation to improve model accuracy
- Leverage transfer learning between models to achieve efficient results with less data and computation
- Build confidence to take on your own project with a modern deep learning framework

Who Should Attend

- Developers
- Algorithm experts
- Software engineers
- Computer vision engineers

Prerequisites

An understanding of fundamental programming concepts in Python such as functions, loops, dictionaries, and arrays.

Course Contents

The Mechanics of Deep Learning
Explore the fundamental mechanics and tools involved in successfully training deep neural networks:

- Train your first computer vision model to learn the process of training.
- Introduce convolutional neural networks to improve accuracy of predictions in vision applications.
- Apply data augmentation to enhance a dataset and improve model generalization.

Pre-trained Models and Recurrent Networks
Leverage pre-trained models to solve deep learning challenges quickly. Train recurrent neural networks on sequential data:

- Integrate a pre-trained image classification model to create an automatic doggy door.
- Leverage transfer learning to create a personalized doggy door that only lets in your dog.
- Train a model to autocomplete text based on New York Times headlines.

Final Project: Object Classification
Apply computer vision to create a model that distinguishes between fresh and rotten fruit.

- Create and train a model that interprets color images.
- Build a data generator to make the most out of small datasets.
- Improve training speed by combining transfer learning and feature extraction.
- Discuss advanced neural network architectures and recent areas of research where students can further improve their skills.

Final Review

- Review key learnings and answer questions.
- Complete the assessment and earn a certificate.
- Complete the workshop survey.
- Learn how to set up your own AI application development environment.