Getting Started with Deep Learning

(NVIDIA)

Duration: 08 hours

Course Prerequisites:

- An understanding of fundamental programming concepts in Python 3 such as functions, loops, dictionaries, and arrays.
- A familiarity of Pandas datastructures.
- An understanding of how to compute a regression line.
- Tools, libraries, frameworks used: TensorFlow 2 with Keras, Pandas

About this Course

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 course, 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.

Learning Objectives

- 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

Topics Covered

- PyTorch
- Convolutional Neural Networks (CNNS)
- Data Augmentation
- Transfer Learning
- Natural Language Processing

Course Outline

Image Classification with the MNIST Dataset	• The MNIST Dataset. • Tensors. • Preparing the Data for Training. • Creating the Model. • Training the Model.
Image Classification of an American Sign Language Dataset	• The American Sign Language Dataset. • Loading the Data. • Visualizing the Data. • Build the Model. • Training the Model.
Convolutional Neural Networks	• Loading and Preparing the Data. • Creating a Convolutional Model. • Summarizing the Model. • Training the Model.
Data Augmentation	• Preparing the Data. • Model Creation. • Data Augmentation. • Training with Augmentation.

	• Saving the Model.
Deploying Your Model	• Loading the Model. • Preparing an Image for the Model. • Making Predictions.
Pre-Trained Models	• An Automated Doggy Door. • Loading the Model. • Loading an Image. • Make a Prediction. • Only Dogs.
Transfer Learning	• Personalized Doggy Door. • Data Augmentation. • The Training Loop. • Fine-Tuning the Model.
Natural Language Processing	• BERT. • Tokenization. • Segmenting Text. • Text Masking. • Question and Answering.