

Accelerating Deep Learning with GPUs

About This Course

Training complex deep learning models with large datasets takes along time. In this course, you will learn how to use accelerated GPU hardware to overcome the scalability problem in deep learning.

You can use accelerated hardware such as Google's Tensor Processing Unit (TPU) or Nvidia GPU to accelerate your convolutional neural network computations time on the Cloud. These chips are specifically designed to support the training of neural networks, as well as the use of trained networks (inference). Accelerated hardware has recently been proven to significantly reduce training time.

But the problem is that your data might be sensitive and you may not feel comfortable uploading it on a public cloud, preferring to analyze it on-premise. In this case, you need to use an in-house system with GPU support. One solution is to use IBM's Power Systems with Nvidia GPU and Power AI. The Power AI platform supports popular machine learning libraries and dependencies including Tensorflow, Caffe, Torch, and Theano.

In this course, you'll understand what GPU-based accelerated hardware is and how it can benefit your deep learning scaling needs. You'll also deploy deep learning networks on GPU accelerated hardware for several problems, including the classification of images and videos.

What will you learn?

- Explain what GPU is, how it can speed up the computation, and its advantages in comparison with CPUs.
- Implement deep learning networks on GPUs.
- Train and deploy deep learning networks for image and video classification as well as for object recognition.

Course Syllabus

- Module 1 - Quick review of Deep Learning
- Module 2 - Hardware Accelerated Deep Learning
- Module 3 - Deep Learning in the Cloud
- Moduel 4 - Distributed Deep Learning