Applied Parallel Computing LLC
http://parallel-computing.pro
+41 78 925 90 90
contact@parallel-computing.pro

4-day Training Course on GPU-enabled Neural Networks

(dates to be arranged)

Deep Learning on GPUs is currently boosting AI's massive leap into real-life applications: autopilots, intelligent automated assistants, real-time translation, image recognition, data sequencing and clustering. With the unprecedented computing power of NVIDIA GPUs, many automotive, robotics and big data companies are creating products and services based on a new class of intelligent machines.

This training course is intended for developers willing to rapidly get NVIDIA-based AI technology into new and existing software solutions. In 4 days we will walk from the necessary theory of Recurrent and Convolutional Neural Networks to practical recipes of using Tensorflow and Caffe frameworks. Finally, course attendees will be guided to build up an AI software stack for a real robot.

Hands-ons: All discussed topics will be accompanied with practical sessions, using Caffe and Tensorflow frameworks with NVIDIA backend. Exercises will be conducted either on attendee's preferred systems, or on the provided GPU server with NVIDIA Tesla K40 and K20 GPUs: https://parallel-computing.pro/test-drive.

All corresponding presentations and code samples will be available to attendees from the beginning of each training day.

Applied Parallel Computing LLC is delivering GPU training courses since 2009. Several dozens of courses have been organized all over Europe, both for commercial and academic customers. We work in close partnership with NVIDIA, CUDA Centers of Excellence and Tesla Preferred Partners. In addition to trainings, our company provides GPU porting/optimization services and CUDA certification.

Day 1: Introduction to Machine learning and Neural Networks

Morning (09:00-12:30)

09:00-10:30: lecture

- An overview of GPU performance in various applications
- Machine learning, image recognition, human recognition abilities
- Image filtering
- Image classification patterns, drawbacks
- Bayes classifier (binary classes), classification tasks, examples
- Regression, approximation
- Neural networks, 13th Hilbert problem
- Artificial neuron, weights, activation functions. Perceptron. Layers. Learning approaches. Backpropagation

10:30-12:30: hands-on session

- hands-on, binary classifier, neural network backpropagation, function approximation problem. (CPU)

Afternoon (13:30-17:00)

13:30-15:00: lecture

- Problems of image classification, objects detection, scene understanding explained
- Overview of natural language processing, speech recognition and translation
- Big data, performance
- CPU vs GPU, NVIDIA Kepler, Maxwell and Pascal architecture & performance
- Automotive: Jetson, Parker, Arduino, Raspberry Pi, Robot Operation System, NVIDIA SDK
- Neural networks, supervised and unsupervised learning
- Testing and training sets. Verification, batches
- Self-Organizing Maps

15:00-17:00: hands-on session

- Hands-on, Kohonen map data clustering, convolutional image filters (CPU)

Day 2: Convolutional Neural Networks and Caffe

Morning (09:00-12:30)

09:00-10:30: lecture

- Caffe framework. Overview, installation
- Blobs, Layers and Nets. Forward and Backward propagation
- Types of layers: Vision, Loss, Activation /Neuron Layers. Data and Common layers
- Model Format. Prototxt, binaryproto - .caffemodel files

10:30-12:30: hands-on session

- Hands-on. MNIST (Handwritten digits recognition), caffe files

Afternoon (13:30-17:00)

13:30-15:00: lecture

- Solvers and methods for neural network training
- Convolutional networks, structures, parameters
- LeNet, AlexNet, GoogleNet, VGG
- Guidelines on choosing neural network structure
- Open-source datasets

15:00-16:30: hands-on session

- Caffe + AlexNet or GoogleNet, different solvers, face recognition (Labelled Faces in the Wild)

Day 3: Recurrent Neural Networks and Tensorflow

Morning (09:00-12:30)

09:00-10:30: lecture

- Overview of Google TensorFlow
- Computation graphs. Building graphs, Sessions
- Tensors (ranks, shapes, types), variables
- tf.train.Optimizer – base class for optimizers
- TensorBoard – suite of visualization tools

10:30-12:30: hands-on session

- TensorBoard. MNIST – handwritten data recognition. Iris dataset

Afternoon (13:30-17:00)

13:30-15:00: lecture

- NVIDIA data flow optimizations and technologies
- Distributed sessions. Creating cluster, distributing training between devices
- Converting data into tensors
- Natural language processing with Recurrent Neural Networks

15:00-16:30: hands-on session

- Hands-on Natural language processing – English-French translator. Principles and limitations

Day 4: NVIDIA DIGITS, TensorRT, DetectNet and final project

Morning (09:00-12:30)

09:00-10:30: lecture

- DIGITS - NVIDIA Deep Learning GPU Training System
- Detectnet – Detecting objects
- NVIDIA TensorRT– optimizing and deploying trained network
- Combining CAFFE, Tensorflow and custom C/C++ code

10:30-12:30: hands-on session

- Recognizing and finding several objects on webcam image

Afternoon (13:30-17:00)

13:30-16:30: hands-on session

- DC motor transient model approximation
- Equalizing and synchronizing DC motors
- Recognizing robot and calculating it's coordinates on image. Detectnet
- Recognizing other objects on scene. Detectnet
- Understanding commands using RNN (Tensorflow)
- Navigating robot, using neural networks for DC motors controlling,
- Avoiding obstacles, understanding and interpreting commands, recognizing objects
- Live demo