Deep Learning and GPU Programming using OpenACC

12 – 14 July 2022

Overview

- The course is co-organised by LRZ, HLRS and NVIDIA Deep Learning Institute (DLI).

- NVIDIA Deep Learning Institute (DLI) offers hands-on training for developers, data scientists, and researchers looking to solve challenging problems with deep learning.

- The workshop combines lectures about Fundamentals of Deep Learning and Fundamentals of Deep Learning for Multi-GPUs with a lecture about Accelerated Computing with OpenACC.

- Learn how to accelerate your applications with OpenACC, how to train and deploy a neural network to solve real-world problems, and how to effectively parallelize training of deep neural networks on Multi-GPUs.

- The lectures are interleaved with many hands-on sessions using Jupyter Notebooks. The exercises will be done on a fully configured GPU-accelerated workstation in the cloud.

DEEP LEARNING INSTITUTE

DLI Mission: Help the world to solve the most challenging problems using AI and deep learning

We help developers, data scientists and engineers to get started in architecting, optimizing, and deploying neural networks to solve real-world problems in diverse industries such as autonomous vehicles, healthcare, robotics, media & entertainment and game development.
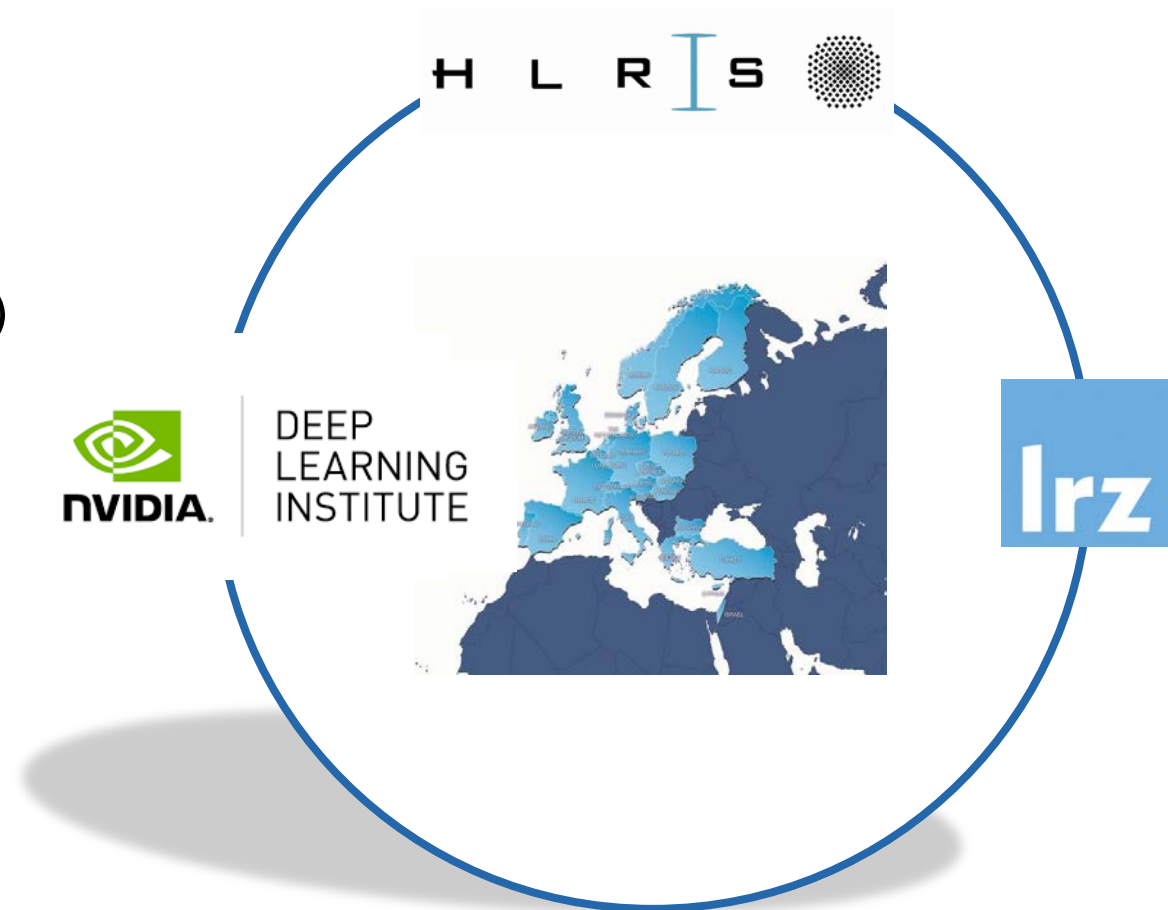
Lecturers

- **Lecturers**:
  - Dr. Momme Allalen (LRZ)
  - PD Dr. Juan Durillo Barrionuevo (LRZ)
  - Dr. Volker Weinberg (LRZ)

- **Teaching Assistants:**
  - Dr. Tobias Haas (HLRS)
  - Dr. Khatuna Kakhiani (HLRS)
  - Dr.-Ing. Lorenzo Zanon (HLRS)
  - Lucienne Dettki (HLRS)

All lecturers are NVIDIA certified University Ambassadors.

1st day: Fundamentals of Accelerated Computing with OpenACC

- On the 1st day you learn the basics of OpenACC, a high-level programming language for programming on GPUs. Discover how to accelerate the performance of your applications beyond the limits of CPU-only programming with simple pragmas.

- You'll learn:
  - How to profile and optimize your CPU-only applications to identify hot spots for acceleration
  - How to use OpenACC directives to GPU accelerate your codebase
  - How to optimize data movement between the CPU and GPU accelerator

- Upon completion, you'll be ready to use OpenACC to GPU accelerate CPU-only applications.

2nd day: Fundamentals of Deep Learning

- Explore the fundamentals of deep learning by training neural networks and using results to improve performance and capabilities.

- During this day, you'll learn the basics of deep learning by training and deploying neural networks. You'll learn how to:
  - Implement common deep learning workflows, such as image classification and object detection
  - Experiment with data, training parameters, network structure, and other strategies to increase performance and capability
  - Deploy your neural networks to start solving real-world problems

- Upon completion, you'll be able to start solving problems on your own with deep learning.

3rd day: Fundamentals of Deep Learning for Multi-GPUs

- The computational requirements of deep neural networks used to enable AI applications like self-driving cars are enormous. A single training cycle can take weeks on a single GPU or even years for larger datasets like those used in self-driving car research. Using multiple GPUs for deep learning can significantly shorten the time required to train lots of data, making solving complex problems with deep learning feasible.

- On the last day we will teach you how to use multiple GPUs to train neural networks. You'll learn:
  - Approaches to multi-GPUs training
  - Algorithmic and engineering challenges to large-scale training
  - Key techniques used to overcome the challenges mentioned above

- Upon completion, you'll be able to effectively parallelize training of deep neural networks using TensorFlow.

Tentative Agenda Day 1: **Fundamentals of Accelerated Computing with OpenACC**

09:00-09:30 Welcome & Intro
09:30-10:30 Profiling
10:30-11:00 Introduction into NVIDIA® Nsight™ Systems
11:00-11:15 Coffee Break
11:15-11:45 Lab 1
11:45-12:15 OpenACC Directives
12:15-12:45 Lab 2
12:45-13:15 GPU Programming

13:15-14:15 Lunch Break

14:15-14:45 Lab 3
14:45-15:15 Data Management
15:15-15:45 Lab 4
15:45-16:00 Coffee Break
16:00-16:30 Loop Optimizations
16:30-16:45 Lab 5
16:45-17:00 Q&A

Tentative Agenda Day 2:
Fundamentals of Deep Learning

9:00-9:20 Welcome and Intro

9:20-11:00 Introduction to Deep Learning and Convolutional Neural Networks

11:00-11:15 Coffee Break

11:00-12:00 How a Neural Network Trains
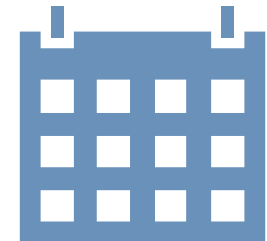
12:00-13:15 Convolutional Neural Networks

13:15-14:15 Lunch Break

14:15-15:45 Data Augmentation, Deployment and Pre-Trained Models

15:45-16:00 Coffee Break

16:00-16:45 Advanced Architectures

16:45-17:00 Q&A

Tentative Agenda Day 3: **Fundamentals of Deep Learning for Multi-GPUs**

9:00-9:15 Intro
9:15-11:00 Stochastic Gradient Descent, a focus session on how training works
11:00-11:15 Coffe Break
11:15-12:00 Hands On Session on Gradient Descent
12:00-13:15 Introduction to Distributed Training
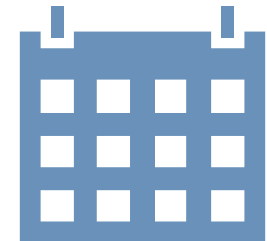
13:15-14:15 Lunch Break

14:15-15:00 Hands On Session on Distributed Training
15:00-15:45 Algorithmic Challenges of Distributed SGD
15:45-16:00 Coffee Break
16:00-16:45 Hands on session on Algorithmic Challenges of Distributed SGD
16:45-17:00 Q&A

Course Webpage

- All slides will be made available during the workshop under:

- https://tinyurl.com/dli-hlrs-2022

- Further information on:

  - Agenda
  - Training Setup
  - Slides
  - Documentation

Training Setup

- To get started, follow these steps:
- Create an NVIDIA Developer account at http://courses.nvidia.com/join Select "Log in with my NVIDIA Account" and then '"Create Account".

- If you use your own laptop, make sure that WebSockets works for you:
  Test your Laptop at http://websocketstest.com
  - Under ENVIRONMENT, confirm that '"WebSockets" is checked yes.
  - Under WEBSOCKETS (PORT 80]. confirm that "Data Receive", "Send", and "Echo Test" are checked yes.
  - If there are issues with WebSockets, try updating your browser.
    We recommend Chrome, Firefox, or Safari for an optimal performance.

- Visit http://courses.nvidia.com/dli-event and enter the event code provided by the instructor.
- You're ready to get started.

And now …

Enjoy the course!