

# Deep Learning and GPU Programming using OpenACC

HLRS | 11 – 13 July 2023

Course material:  
<https://tinyurl.com/dl-openacc-2023>



- This course is organised by HLRS in cooperation with LRZ.
- Previously a similar course was offered within the NVIDIA Deep Learning Institute (DLI) University Ambassador programme, which has been suspended in February 2023. The material has been heavily revised and in part completely newly developed for this course.
- The workshop combines an introduction to Deep Learning and Deep Learning for Multi-GPUs with a lecture on 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 parallelise training of deep neural networks on Multi-GPUs.
- The lectures are interleaved with many hands-on sessions partly using Jupyter Notebooks. The exercises will be done on the AI partition of HLRS's cluster HAWK.

# Lecturers




- **Lecturers:**

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



All lecturers are NVIDIA certified University Ambassadors.

# 1<sup>st</sup> day: Fundamentals of Accelerated Computing with OpenACC




- On the 1<sup>st</sup> 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 optimise 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.

## 2<sup>nd</sup> day: Introduction to 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.

# 3<sup>rd</sup> day: Introduction to Deep Learning for Multi-GPUs



- 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 parallelise training of deep neural networks using TensorFlow.

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



09:00-09:30 Welcome & Intro

09:30-11:00 Profiling

**11:00-11:15 Coffee Break**

11:15-12:15 Introduction into NVIDIA® Nsight™ Systems

12:15-12:45 Lab 1

12:45-13:15 OpenACC Directives

**13:15-14:30 Lunch**

14:30-14:45 Lab 2

14:45-15:15 GPU Programming


15:15-15:45 Lab 3

**15:45-16:00 Coffee Break**


16:00-16:30 Data Management

16:30-16:45 Lab 4

16:45-17:00 Q&A, Final Remarks



# Core Agenda All Days



09:00-11:00 Part 1

**11:00-11:15 Coffee Break**

11:00-13:15 Part 2


**13:15-14:30 Lunch Break**

14:30-15:45 Part 3

**15:45-16:00 Coffee Break**


16:00-16:45 Part 4

16:45-17:00 Q&A and Final Remarks





# Tentative Agenda Day 2: Introduction to Deep Learning



09:00-09:20 Welcome and Intro

09:20-11:00 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


**13:15-14:30 Lunch Break**

14:30-15:45 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: Introduction to Deep Learning for Multi-GPUs



09:00-09:15 Intro

09:15-11:00 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

**13:15-14:30 Lunch Break**


14:30-15:00 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/dl-openacc-2023>

- Further information on:

- Agenda
- Training Setup
- Slides
- Documentation



And now ...



**Enjoy the course!**