

# Data Parallelism - How to Train Deep Learning Models on Multiple GPUs

LRZ | 5 October 2023

Course material:  
<https://tinyurl.com/hdli1w23>



# Overview



- The course is co-organised by LRZ 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.
- This course teaches you techniques for data-parallel deep learning training on multiple GPUs to shorten the training time required for data-intensive applications. Working with deep learning tools, frameworks, and workflows to perform neural network training, you'll learn how to decrease model training time by distributing data to multiple GPUs, while retaining the accuracy of training on a single GPU.
- 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



DEEP  
LEARNING  
INSTITUTE



- **Lecturer:**

- PD Dr. Juan Durillo Barrionuevo (LRZ)

All lecturers of our DLI courses are NVIDIA certified University Ambassadors.



# Data Parallelism - How to Train Deep Learning Models on Multiple GPUs



- By participating in this course, you'll:
  - Understand how data parallel deep learning training is performed using multiple GPUs
  - Achieve maximum throughput when training, for the best use of multiple GPUs
  - Distribute training to multiple GPUs using Pytorch Distributed Data Parallel
  - Understand and utilize algorithmic considerations specific to multi-GPU training performance and accuracy

# Tentative Agenda



10:00-10:15 Introduction

10:15-11:15 Neural Network Training and Stochastic Gradient Descent

**11:15-11:30 Coffee Break**

11:30-12:30 Neural Network Training and Intro to Parallel Training


**12:30-13:30 Lunch Break**

13:30-15:00 Data Parallelism using Pytorch Distributed Data Parallel

**15:00-15:15 Coffee break**

15:15-16:45 Challenges of Data Parallel using Multiple GPUs

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



# Course Webpage



DEEP  
LEARNING  
INSTITUTE



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

- <https://tinyurl.com/hdli1w23>



- Further information on:

- Agenda
- Training Setup
- Slides
- Documentation



And now ...



DEEP  
LEARNING  
INSTITUTE



**Enjoy the course!**