



DEEP  
LEARNING  
INSTITUTE

# Deep Learning and GPU Programming Workshop

9 – 12 September 2024

Workshop material:  
<https://tinyurl.com/hdli2s24>



# Overview



- The workshop 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.
- The online workshop combines lectures about Accelerated Computing with OpenACC and CUDA on single and multiple GPUs with lectures about Fundamentals of Deep Learning.
- Learn how to accelerate your applications with OpenACC and CUDA and how to train and deploy a neural network to solve real-world problems.
- 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



## 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



- **Lecturers:**

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



DEEP  
LEARNING  
INSTITUTE



All instructors are NVIDIA certified University Ambassadors.

1<sup>st</sup> day:

# Fundamentals of Accelerated Computing with OpenACC



DEEP  
LEARNING  
INSTITUTE



- Learning Objectives:
  - Profile and optimize your CPU-only applications to identify hot spots for acceleration.
  - Use OpenACC directives to GPU-accelerate your codebase.
  - Optimize data movement between the CPU and GPU accelerator.

2<sup>nd</sup> day:

## Fundamentals of Accelerated Computing with CUDA C/C++



DEEP  
LEARNING  
INSTITUTE



- At the conclusion of the workshop, you'll have an understanding of the fundamental tools and techniques for GPU-accelerating C/C++ applications with CUDA and be able to:
  - Write code to be executed by a GPU accelerator
  - Expose and express data and instruction-level parallelism in C/C++ applications using CUDA
  - Utilize CUDA-managed memory and optimize memory migration using asynchronous prefetching
  - Leverage command-line and visual profilers to guide your work
  - Utilize concurrent streams for instruction-level parallelism
  - Write GPU-accelerated CUDA C/C++ applications, or refactor existing CPU-only applications, using a profile-driven approach



- Learning Objectives

- Use concurrent CUDA streams to overlap memory transfers with GPU computation
- Utilize all available GPUs on a single node to scale workloads across all available GPUs
- Combine the use of copy/compute overlap with multiple GPUs
- Rely on the NVIDIA Nsight™ Systems Visual Profiler timeline to observe improvement opportunities and the impact of the techniques covered in the workshop

# 4<sup>th</sup> day: Fundamentals of Deep Learning



- Learning Objectives:
  - Learn the fundamental techniques and tools required to train a deep learning model
  - Gain experience with common deep learning data types and model architectures
  - Enhance datasets through data augmentation to improve model accuracy
  - Leverage transfer learning between models to achieve efficient results with less data and computation
  - Build confidence to take on your own project with a modern deep learning framework



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



10:00-10:15 Intro

10:15-12:00 Profiling


**12:00-13:00 Lunch Break**

13:00-14:20 OpenACC Directives

**14:20-14:30 Coffee Break**

14:30-15:45 GPU Programming and Data Management

15:45-16:00 Q&A, Final Remarks



**All times are in CEST**


# Tentative Agenda Day 2: Fundamentals of Accelerated Computing with CUDA C/C++



10:00-10:15 Introduction CUDA C/C++

10:15-12:00 Accelerating Applications with CUDA C/C++

**12:00-13:00 Lunch Break**



**All times are in CEST**

13:00-14:20 Managing Accelerated Application Memory with CUDA unified memory and nsys

**14:20-14:30 Coffee Break**

14:30-15:45 Asynchronous Streaming and Visual Profiling for Accelerated Applications with CUDA C/C++

15:45-16:00 Q&A, Final Remarks


# Tentative Agenda Day 3: Accelerating CUDA C++ Applications with Multiple GPUs



10:00-10:15 Welcome and Intro, Tools Overview

10:15-12:00 Introduction & Main Objectives

**12:00-13:00 Lunch Break**



**All times are in CEST**

13:00-14:20 Copy/Compute Overlap: Kernel Launches and Memory Copies in Non-Default Streams

**14:20-14:30 Coffee Break**

14:30-15:45 Multiple GPUs

15:45-16:00 Q&A, Final Remarks


# Tentative Agenda Day 4: Fundamentals of Deep Learning



10:00-10:20 Welcome and Intro

10:20-12:00 Introduction to Deep Learning and Convolutional Neural Networks

**12:00-13:00 Lunch Break**



**All times are in CEST**

13:00-14:20 Data Augmentation, Deployment and Pre-Trained Models

**14:20-14:30 Coffee Break**

14:30-15:45 Advanced Architectures

15:45-16:00 Q&A

# Workshop Webpage



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

- <https://tinyurl.com/hdli2s24>



- Further information on:

- Agenda
- Training Setup
- Slides
- Documentation



# Training Setup




- To get started, follow these steps:
- Create an NVIDIA Developer account at <https://learn.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 or Firefox for an optimal performance.
- Visit <https://learn.nvidia.com/dli-event> and enter the event code provided by the instructor.
- You're ready to get started.

# Course Datasheets by NVIDIA

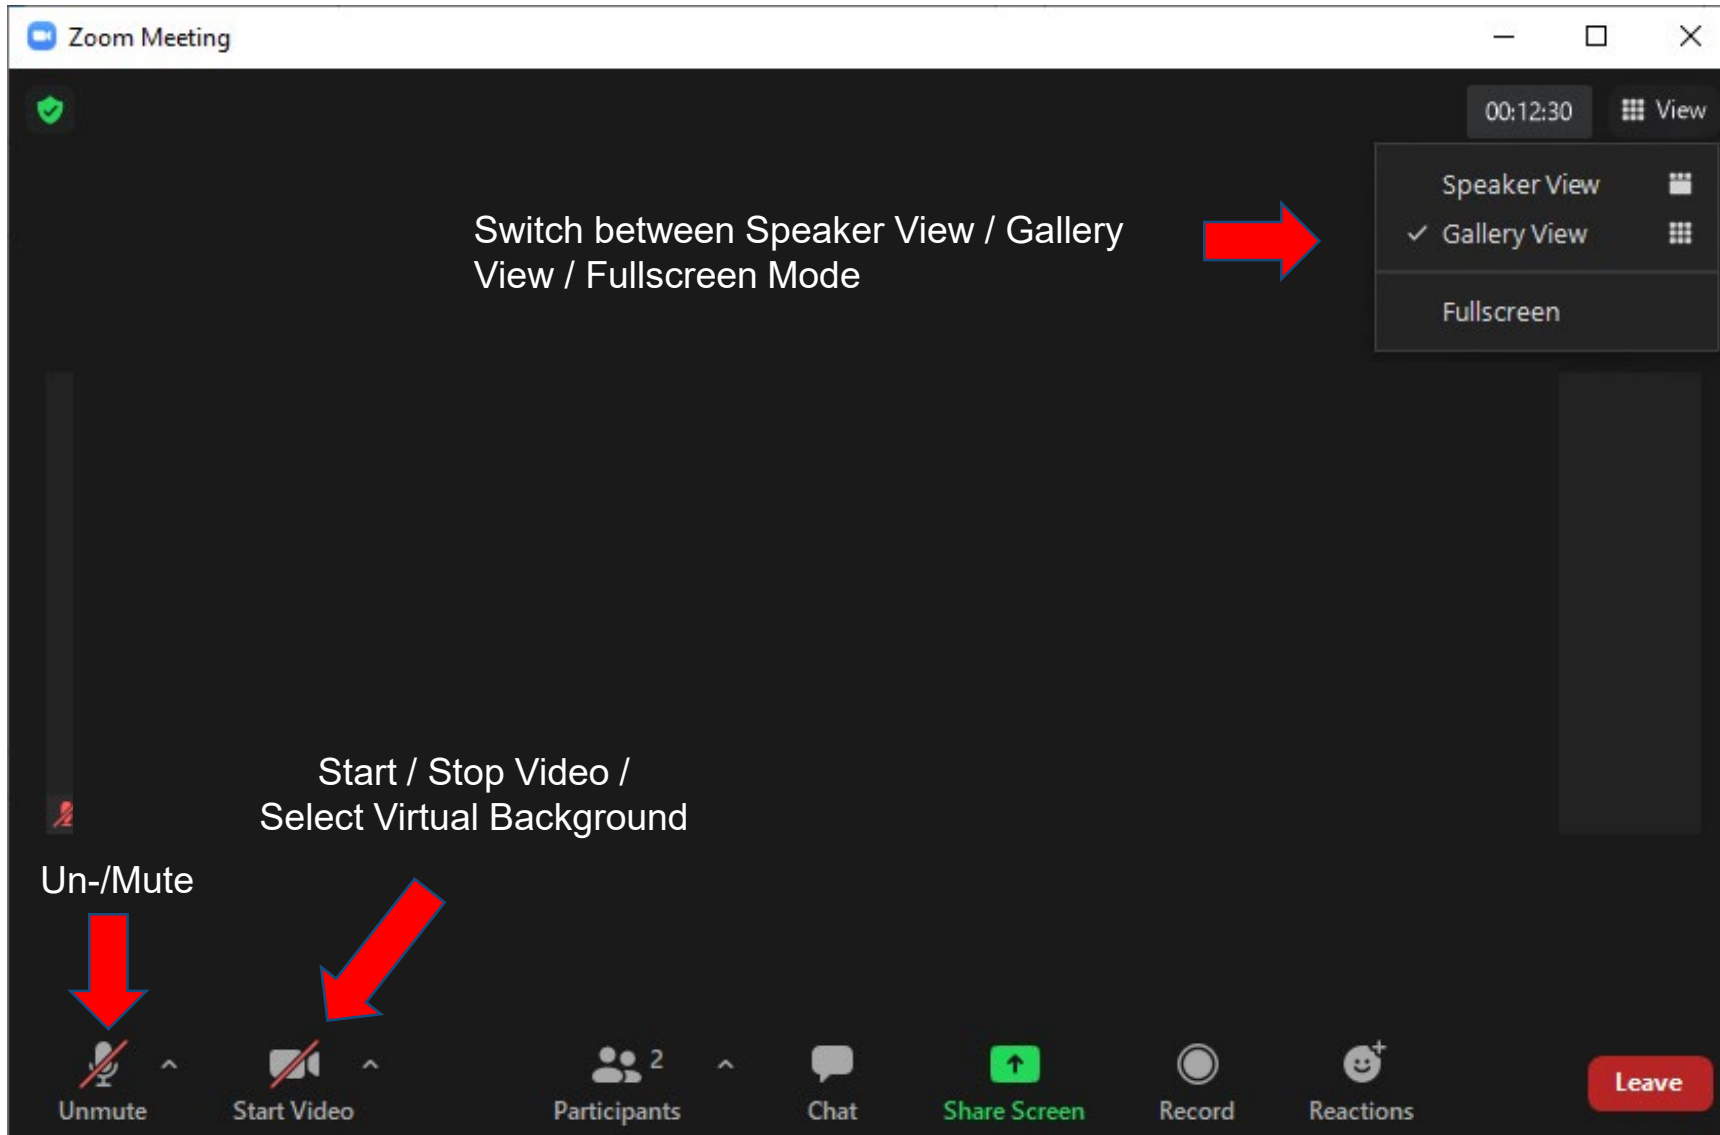


- <https://developer.nvidia.com/dli/getready>
- **Prepare For Your NVIDIA DLI Training**
- **Course Datasheets**
  - Please review the workshop datasheet, which includes prerequisites, agenda, suggested material, and resources for continued learning.



- To ensure a pleasant experience with Zoom Meeting, we encourage participants to **download and install the latest Zoom application** via <https://zoom.us/download>.
- If you have problems with your computer audio, you can also **join by phone**.  
Find your local number: <https://lrz-de.zoom-x.de/u/cpBvnN2sD>





Zoom Meeting

00:12:30 View

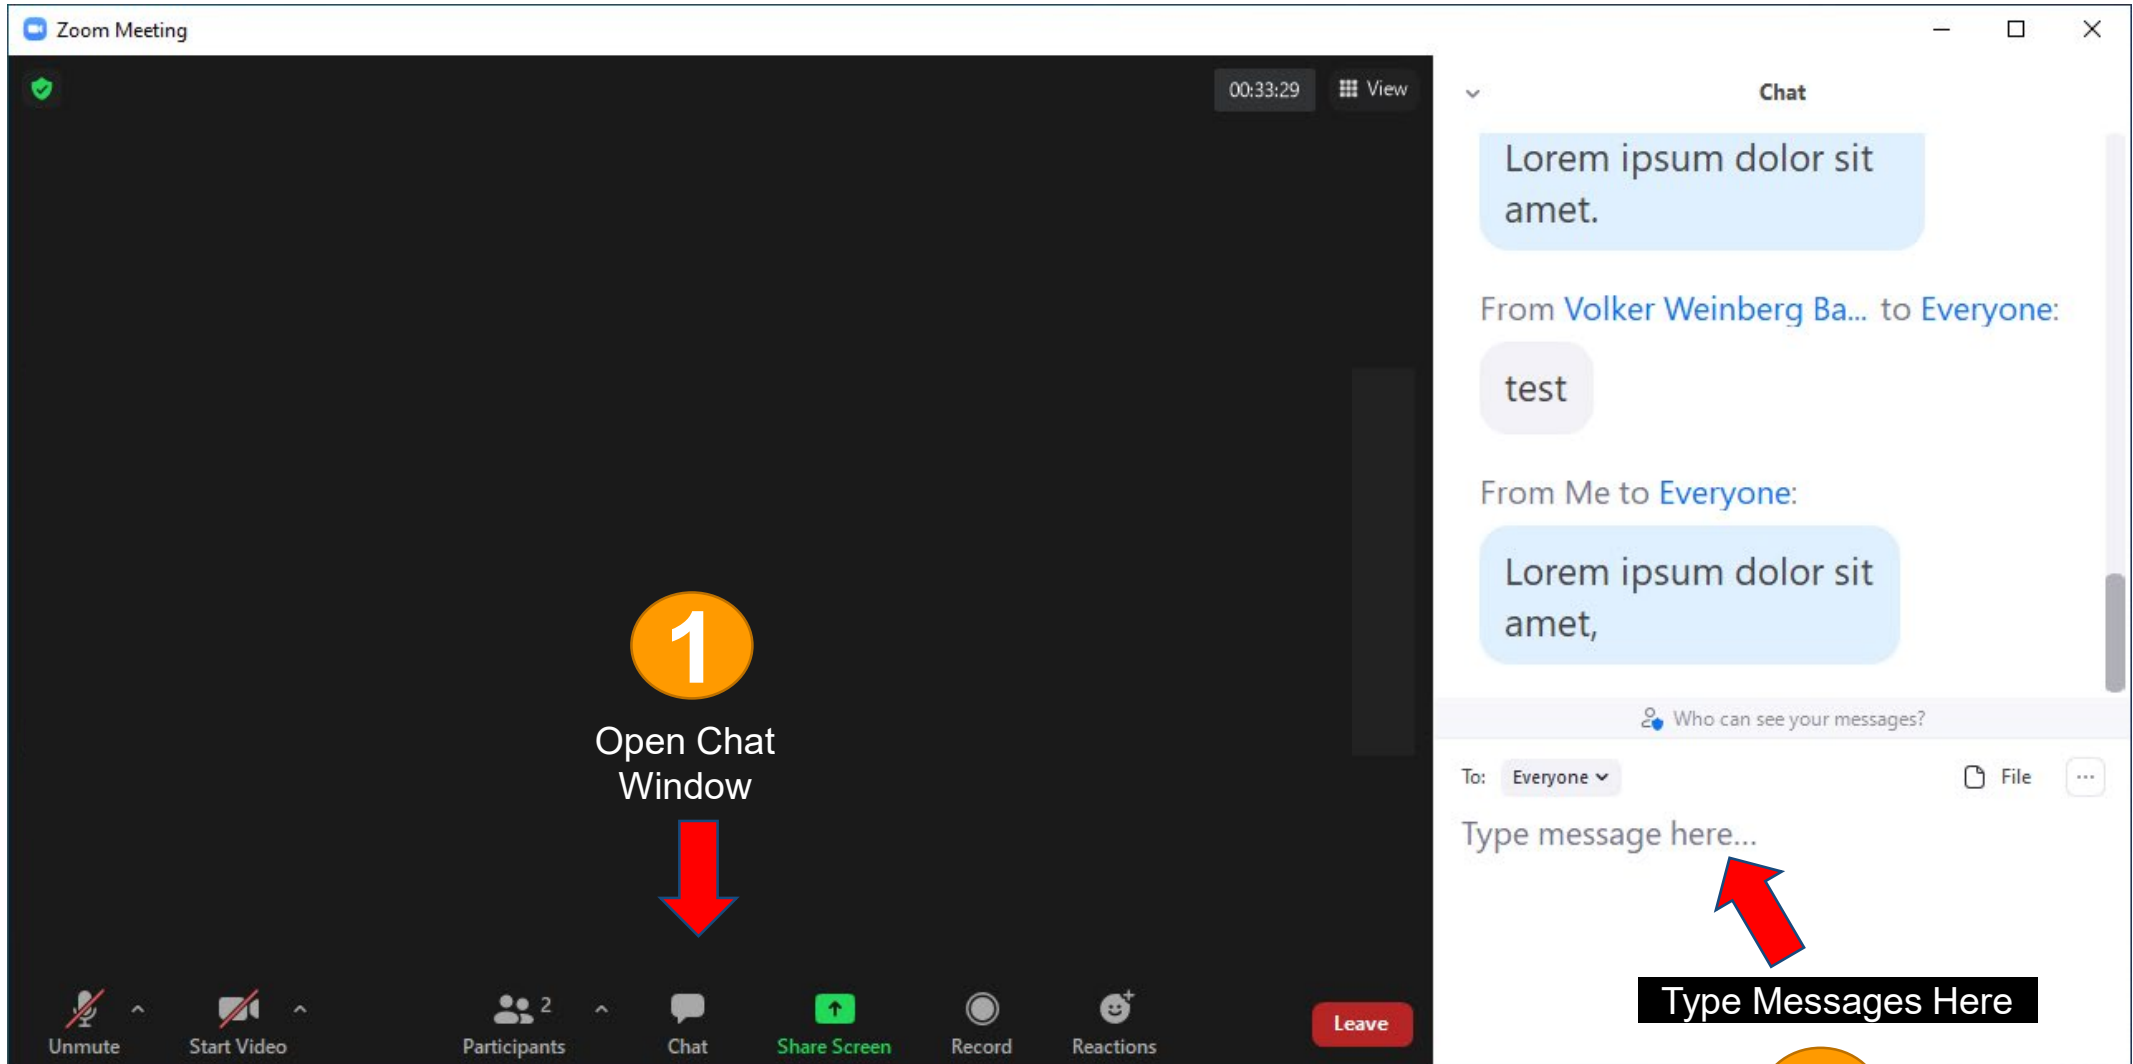
Switch between Speaker View / Gallery View / Fullscreen Mode

Un-/Mute

Start / Stop Video / Select Virtual Background

Unmute Start Video Participants Chat Share Screen Record Reactions Leave

The screenshot shows the Zoom meeting interface with several annotations. A red arrow points from the text 'Switch between Speaker View / Gallery View / Fullscreen Mode' to the 'View' dropdown menu, which is open and shows 'Speaker View', 'Gallery View' (checked), and 'Fullscreen'. Another red arrow points from the text 'Un-/Mute' to the 'Unmute' button. A third red arrow points from the text 'Start / Stop Video / Select Virtual Background' to the 'Start Video' button. The bottom toolbar includes buttons for 'Unmute', 'Start Video', 'Participants', 'Chat', 'Share Screen', 'Record', 'Reactions', and 'Leave'.



Zoom Meeting

00:33:29 View

Chat

Lorem ipsum dolor sit amet.

From Volker Weinberg Ba... to Everyone:

test

From Me to Everyone:

Lorem ipsum dolor sit amet,

Who can see your messages?

To: Everyone

Type message here...

File

Unmute Start Video Participants Chat Share Screen Record Reactions Leave

1

Open Chat Window

2

Type Messages Here

# zoom

## Participants List



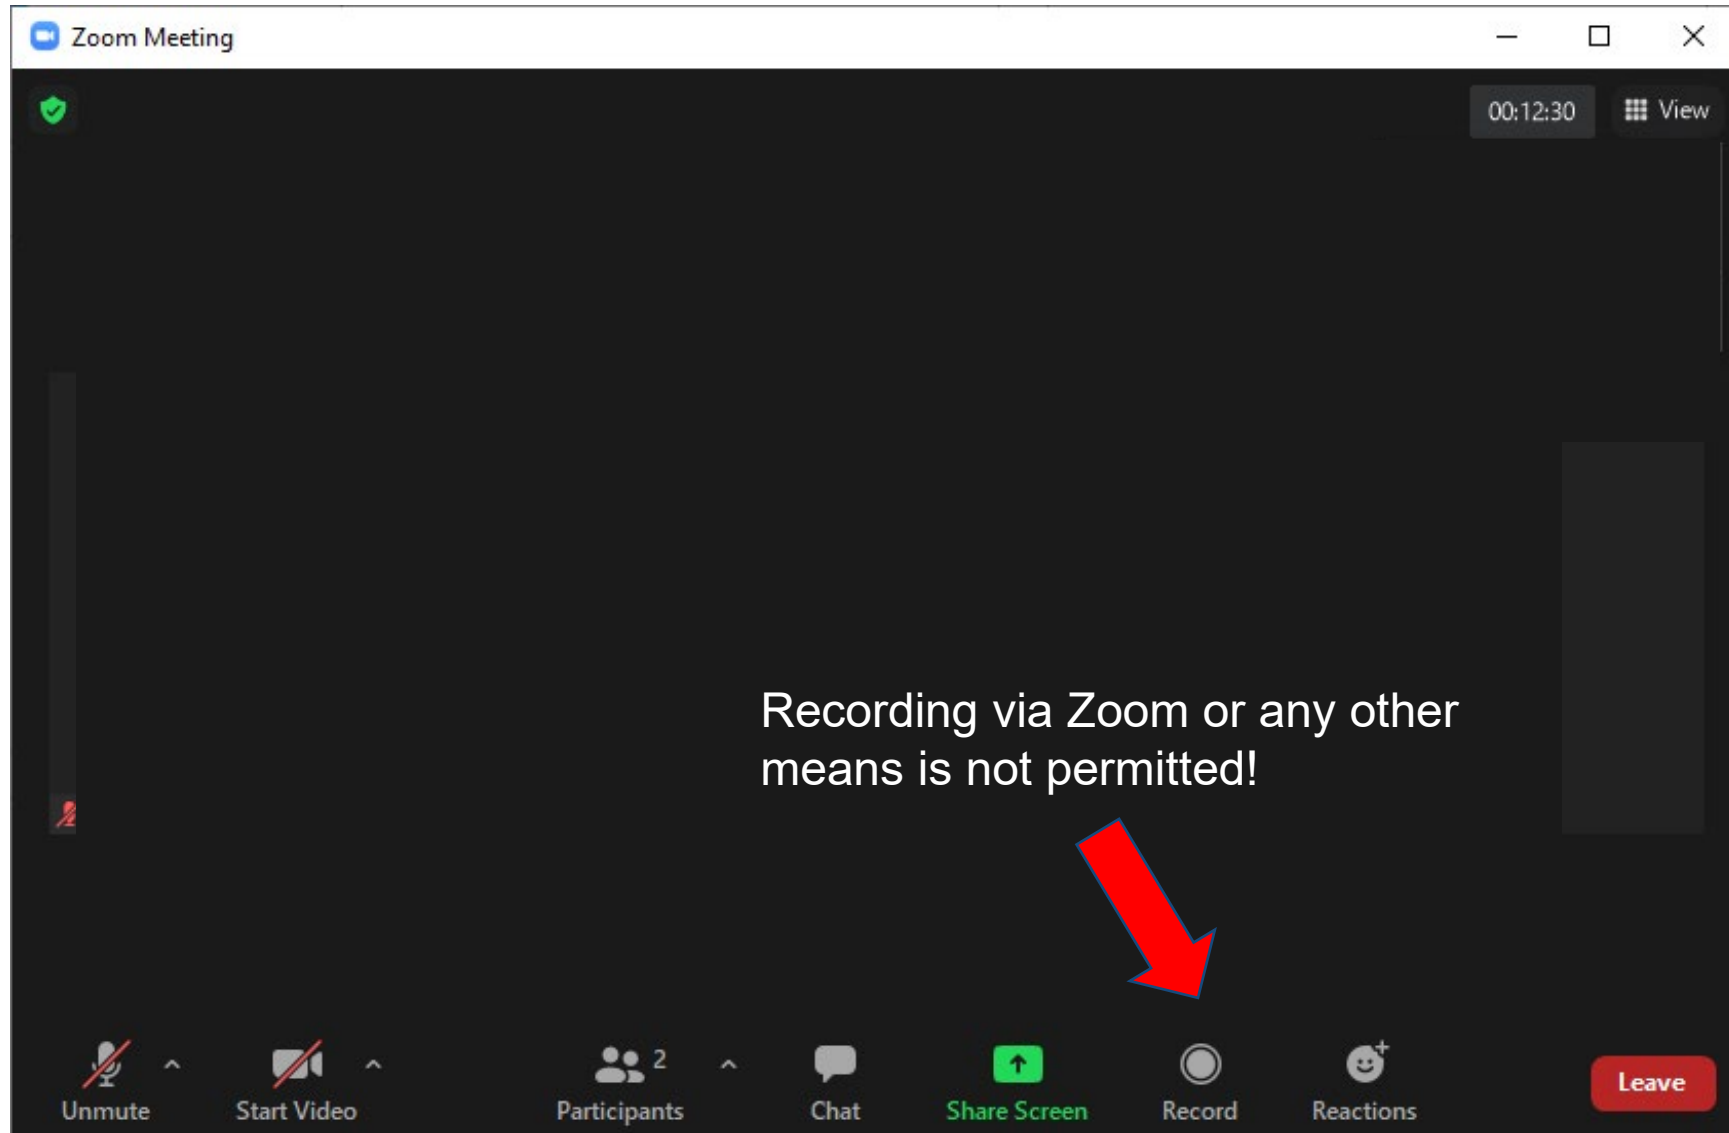
Kindly use “<first name> <last name> (<institute>)” as your screenname.

Otherwise you will not receive a certificate of attendance after the course.


The screenshot shows a Zoom Meeting window with a dark theme. The top bar displays 'Zoom Meeting', a green checkmark, a timer at '00:19:17', and a 'View' button. The main area is mostly black. At the bottom, there is a toolbar with icons for Unmute, Start Video, Participants (with a '2' next to it), Chat, Share Screen, Record, Reactions, and a red 'Leave' button. On the right side, a 'Participants (2)' panel is open, showing two participants: 'Volker Wei... (Me)' and 'Volker Weinberg Backup (Host)'. The 'Me' participant has 'Unmute' and 'More >' buttons. A blue-bordered context menu is open over the 'More >' button, containing 'Edit Profile Picture' and 'Rename' options. A red arrow points from the 'Rename Yourself' text below to the 'Rename' option in the menu. A large orange circle with the number '1' is positioned over the 'Participants' icon in the toolbar, with the text 'Open Participants List' and a red arrow pointing down to the 'Participants' icon. Another large orange circle with the number '2' is positioned below the 'Rename Yourself' text.

1  
Open Participants List

2  
Rename Yourself



- Please **raise your hand** if you have questions (of general interest).
- You can also use **chat window** to ask questions.
- If you do not mind, please **show your video when asking questions** to make this course as interactive as possible.
- **Push to Talk:** The Push to Talk feature allows you to remain muted throughout the Zoom meeting and only if you hold down the spacebar you will be unmuted.
- **Instant Feedback:**



And now ...



DEEP  
LEARNING  
INSTITUTE



**Enjoy the workshop!**