





DEEP  
LEARNING  
INSTITUTE



# Deep Learning and GPU programming using OpenACC

14 – 17 July 2020



# Overview



- The course combines lectures about *Fundamentals of Deep Learning for Computer Vision and Multiple Data Types* with lectures about *Accelerated Computing with OpenACC and DL on HLRS systems*.
- The first three days are organized by LRZ and NVIDIA. All instructors are NVIDIA certified University Ambassadors. On the last day, you will learn more about data preparation and DL on the systems at HLRS.
- Learn how to train and deploy a neural network to solve real-world problems, how to generate effective descriptions of content within images and video clips, and how to accelerate your applications with OpenACC.
- The lectures are interleaved with many hands-on sessions using Jupyter Notebooks. The exercises on the first 3 days will be done on a fully configured GPU-accelerated workstation in the cloud, on the last day on an HLRS system.



## 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

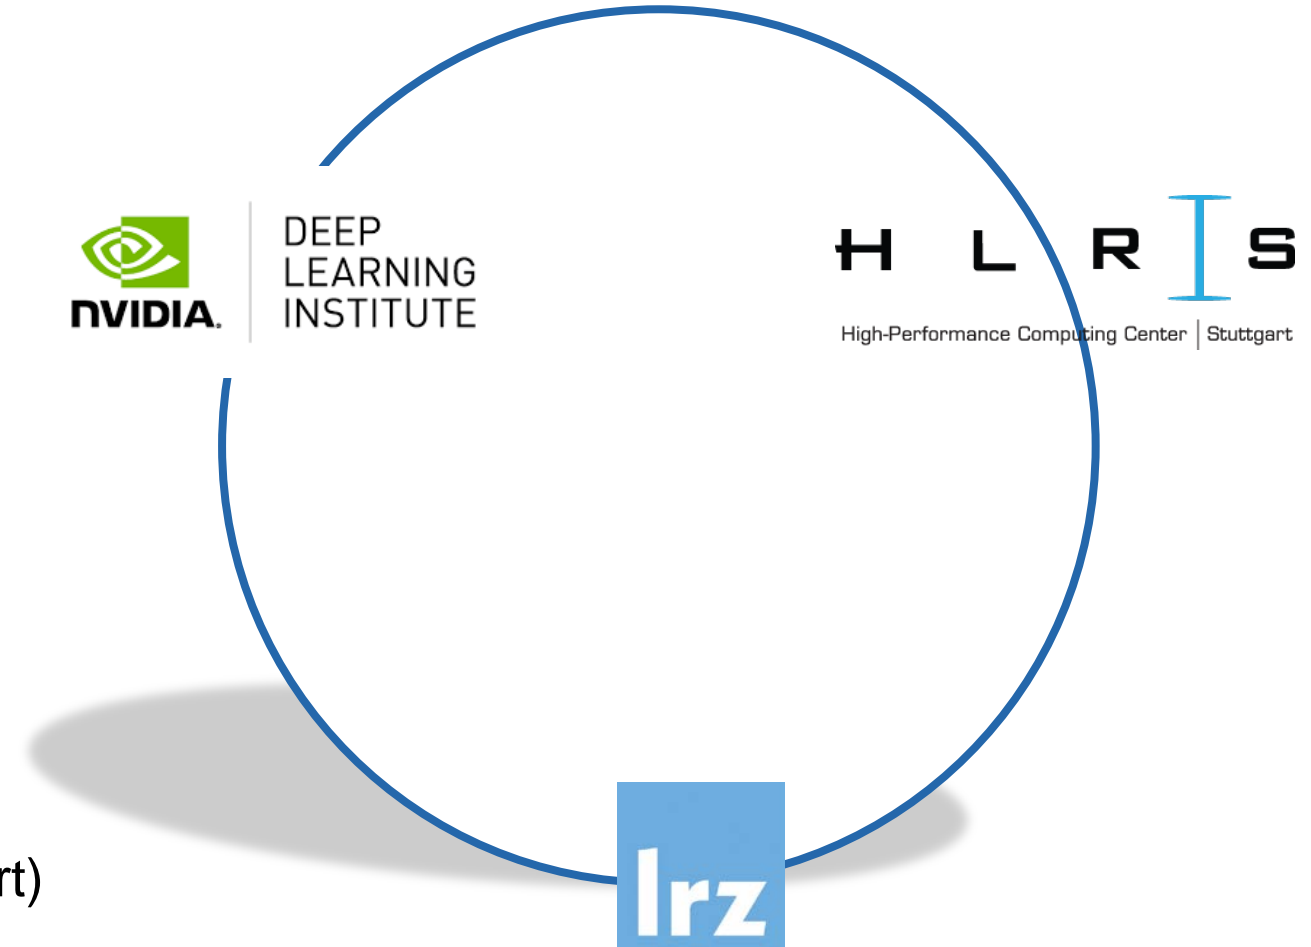


- **Lecturers:**

- Dr. Momme Allalen (LRZ)
- PD Dr. Juan Durillo Barrionuevo (LRZ)
- Dr. Volker Weinberg (LRZ)
  
- Dr. Lorenzo Zanon (HLRS)
- Patrick Vogler, HLRS in cooperation with EXCELLERAT
- Dr. Khatuna Kakhiani (HLRS)

- **Invited Speaker:**

- Dr.-Ing. Andrea Beck (University of Stuttgart)



# 1<sup>st</sup> day: Fundamentals of Deep Learning for Computer Vision

- 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.

# 2<sup>nd</sup> day: Fundamentals of Deep Learning for Multiple Data Types



- The 2<sup>nd</sup> day explores how convolutional and recurrent neural networks can be combined to generate effective descriptions of content within images and video clips.
- Learn how to train a network using TensorFlow and the Microsoft Common Objects in Context (COCO) dataset to generate captions from images and video by:
  - Implementing deep learning workflows like image segmentation and text generation
  - Comparing and contrasting data types, workflows, and frameworks
  - Combining computer vision and natural language processing
- Upon completion, you'll be able to solve deep learning problems that require multiple types of data inputs.

# 3<sup>rd</sup> day: Guest Lecture + Fundamentals of Accelerated Computing with OpenACC



- On the 3<sup>rd</sup> day we start with a **guest lecture** about **Deep Neural Networks for Data-Driven Turbulence Models**, using DL in CFD by Dr.-Ing. Andrea Beck, Institute of Aerodynamics and Gas Dynamics, University of Stuttgart
- Then 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.



# 4<sup>th</sup> day: ML Examples and Methods on HLRS systems




- On the 4<sup>th</sup> day we cover Pre-processing of data, mathematical methods and machine learning as well as Deep learning on HLRS Systems.
- Focus on Pre-processing, Feature Engineering and Machine Learning: Stuttgart S-Bahn Example ([Dr. Lorenzo Zanon](#), Li Zhong and Dennis Hoppe, HLRS)
- Data Compression of numerical data sets with the BigWhoop library (Patrick Vogler, HLRS in cooperation with EXCELLERAT)
- Waste classification using deep learning (Dr. Khatuna Kakhiani)



# Tentative Agenda Day 1: Fundamentals of Deep Learning for Computer Vision

- 09:00-09:20 Welcome and Intro
- 09:20-11:00 Introduction to Deep Learning and Training (I)
- **11:00-11:15 Coffee Break**
- 11:15-12:00 Introduction to Deep Learning and Training (II)
  
- **12:00-13:00 Lunch Break**
  
- 13:00-15:00 Deploying Neural Networks and Performance
- **15:00-15:15 Coffee Break**
- **15:15 Group Picture**
- 15:15-16:45 Beyond classification: Object Recognition
- 16:45-17:00 Q&A



# Tentative Agenda Day 2: Fundamentals of Deep Learning for Multiple Data Types

- 09:00-09:20 Welcome and Intro
- 09:20-12:00 Introduction to CNNs and Object Segmentation (I)
- **11:00-11:15 Coffee Break**
- 11:15-12:00 Introduction to CNNs and Object Segmentation (II)
  
- **12:00-13:00 Lunch Break**
  
- 13:00-15:00 Word Generation with RNNs
- **15:00-15:15 Coffee Break**
- 15:15-16:45 Image Captioning by Combining RNNs and CNNs
- 16:45-17:00 Q&A



# Tentative Agenda Day 3: **Fundamentals of Accelerated Computing with OpenACC**



- 09:00-10:30 Guest Lecture *Dr.-Ing. Andrea Beck*
- 10:30-10:45 **Coffee Break**
- 10:45-11:00 Intro
- 11:00-12:00 Profiling
  
- **12:00-13:00 Lunch Break**
  
- 13:00-15:00 OpenACC Directives
- **15:00-15:15 Coffee Break**
- 15:00-16:00 GPU Programming
- 16:00-16:45 Data Management and Loop Optimizations
- 16:45-17:00 Q&A, Final Remarks



# Tentative Agenda Day 4: **ML Examples and Methods on HLRS systems**

- 09:00-12:30 Focus on Pre-processing, Feature Engineering and Machine Learning: Stuttgart S-Bahn Example (Dr. Lorenzo Zanon - presenter, Li Zhong and Dennis Hoppe, HLRS)
- **12:30-13:30 Lunch Break**
- 13:30-14:30 Data Compression of numerical data sets with the BigWhoop library (Patrick Vogler, HLRS in cooperation with EXCELLERAT)
- 14:30-16:30 Waste classification using Deep learning (Dr. Khatuna Kakhiani, HLRS)



# Workshop Webpage

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

- <https://tinyurl.com/dl-openacc>



- Further information on:

- Agenda
- Training Setup
- Documentation
- ZOOM



# 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.

# zoom

## Joining the Meeting

### Join Zoom Meeting

<https://us02web.zoom.us/j/82140544180?pwd=ZVRYejBDVFd2TTRDOTdiSjJ4K1hscz09>

Meeting ID: 821 4054 4180

Password: 9wDTcu

If you have problems with your computer audio, you can also join by phone.



### Dial by your location

+49 30 5679 5800 Germany


+49 695 050 2596 Germany

+49 69 7104 9922 Germany

Meeting ID: 821 4054 4180

Password: 774450

Find your local number: <https://us02web.zoom.us/u/kQA2T45TJ>

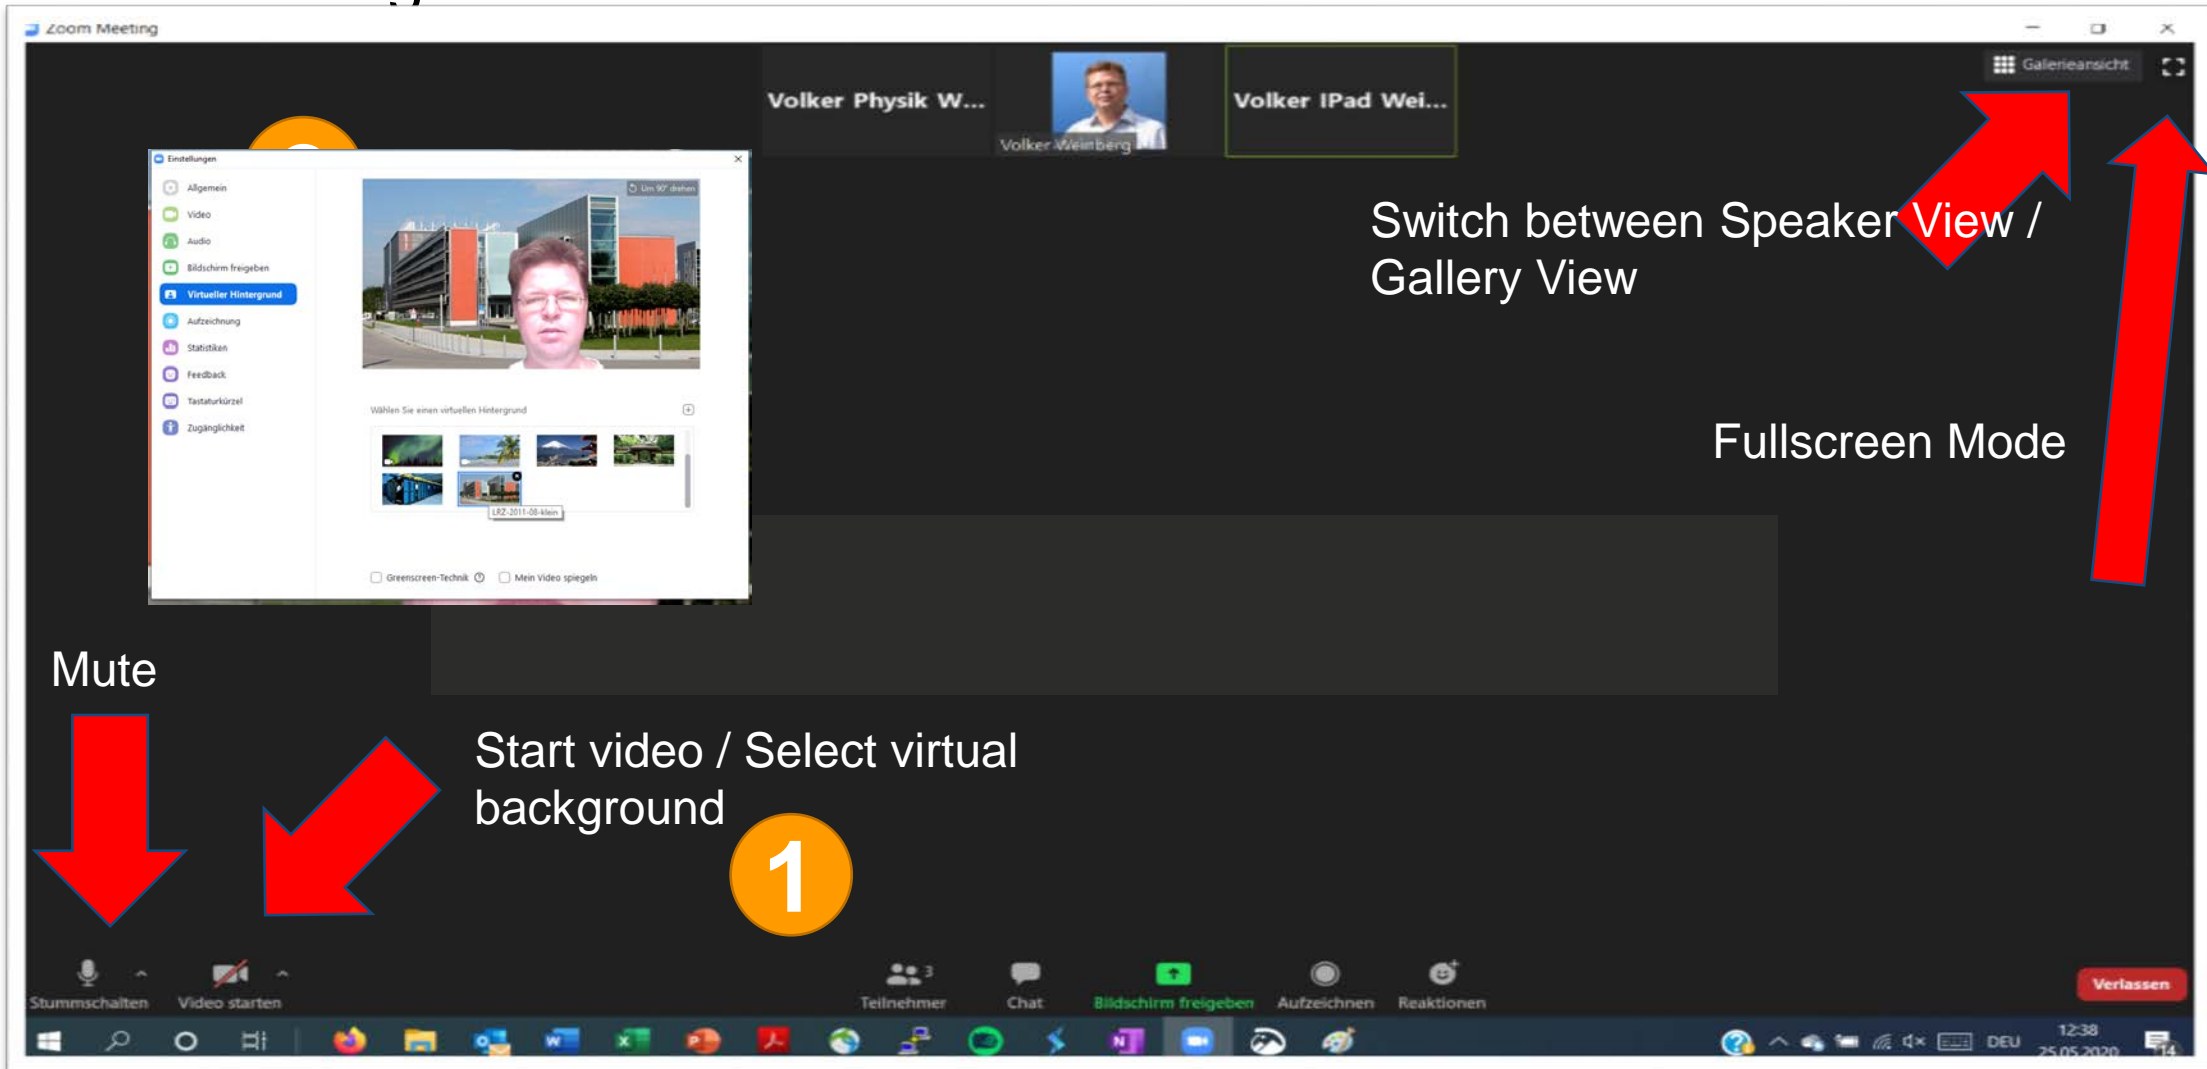




- ZOOM help centre has great resources with help articles and videos for getting started: <https://support.zoom.us/hc/en-us>.
- This “Getting Started” page is a great resource: <https://support.zoom.us/hc/en-us/categories/200101697>
- This FAQ has tons of useful info: <https://support.zoom.us/hc/en-us/articles/206175806-Frequently-Asked-Questions>.
- We strongly encourage you to read some of the basic info relevant to your operating system:
  - Getting Started on Windows and Mac: <https://support.zoom.us/hc/en-us/articles/201362033-Getting-Started-on-Windows-and-Mac>.
  - Getting Started On Chrome OS: <https://support.zoom.us/hc/en-us/articles/213298746-Getting-Started-On-Chrome-OS>.
- You may log in via the app or a browser. We recommend downloading the app for the best experience. It may take several minutes to download, so if you are using Zoom for the first time, please download the app prior to the event, <https://zoom.us/download> , or join the meeting early.
- You may sign up for a free account at [zoom.us/signup](https://zoom.us/signup). Or, you may join a meeting as a guest without a Zoom account.



# Speaker View vs. Gallery View, Virtual Background



The screenshot shows a Zoom meeting window with a dark theme. At the top, there are three participant thumbnails: 'Volker Physik W...', 'Volker Weinberg' (with a video feed), and 'Volker iPad Wei...'. A green box highlights the 'Volker iPad Wei...' thumbnail. On the right side, a 'Teilnehmer (3)' panel lists the participants: 'Volker Physik Weinberg (Ich)', 'Volker Weinberg (Host)', and 'Volker iPad Weinberg'. Below this panel is a reaction toolbar with icons for thumbs down, thumbs up, clapping hands, a coffee cup, and a clock. A tooltip labeled 'Weitere nonverbale Rückmeldungen' is shown over the reaction icons. At the bottom of the meeting window, there is a control bar with buttons for 'Stummschalten', 'Video starten', 'Teilnehmer' (with a '3' icon), 'Chat', 'Bildschirm freigeben', 'Aufzeichnen', 'Reaktionen', and 'Verlassen'. A red arrow labeled '1' points to the 'Teilnehmer' button. On the right side, a zoomed-in view of the reaction toolbar is shown, with a red arrow labeled '2' pointing to the 'Reaktionen' button in the bottom control bar.

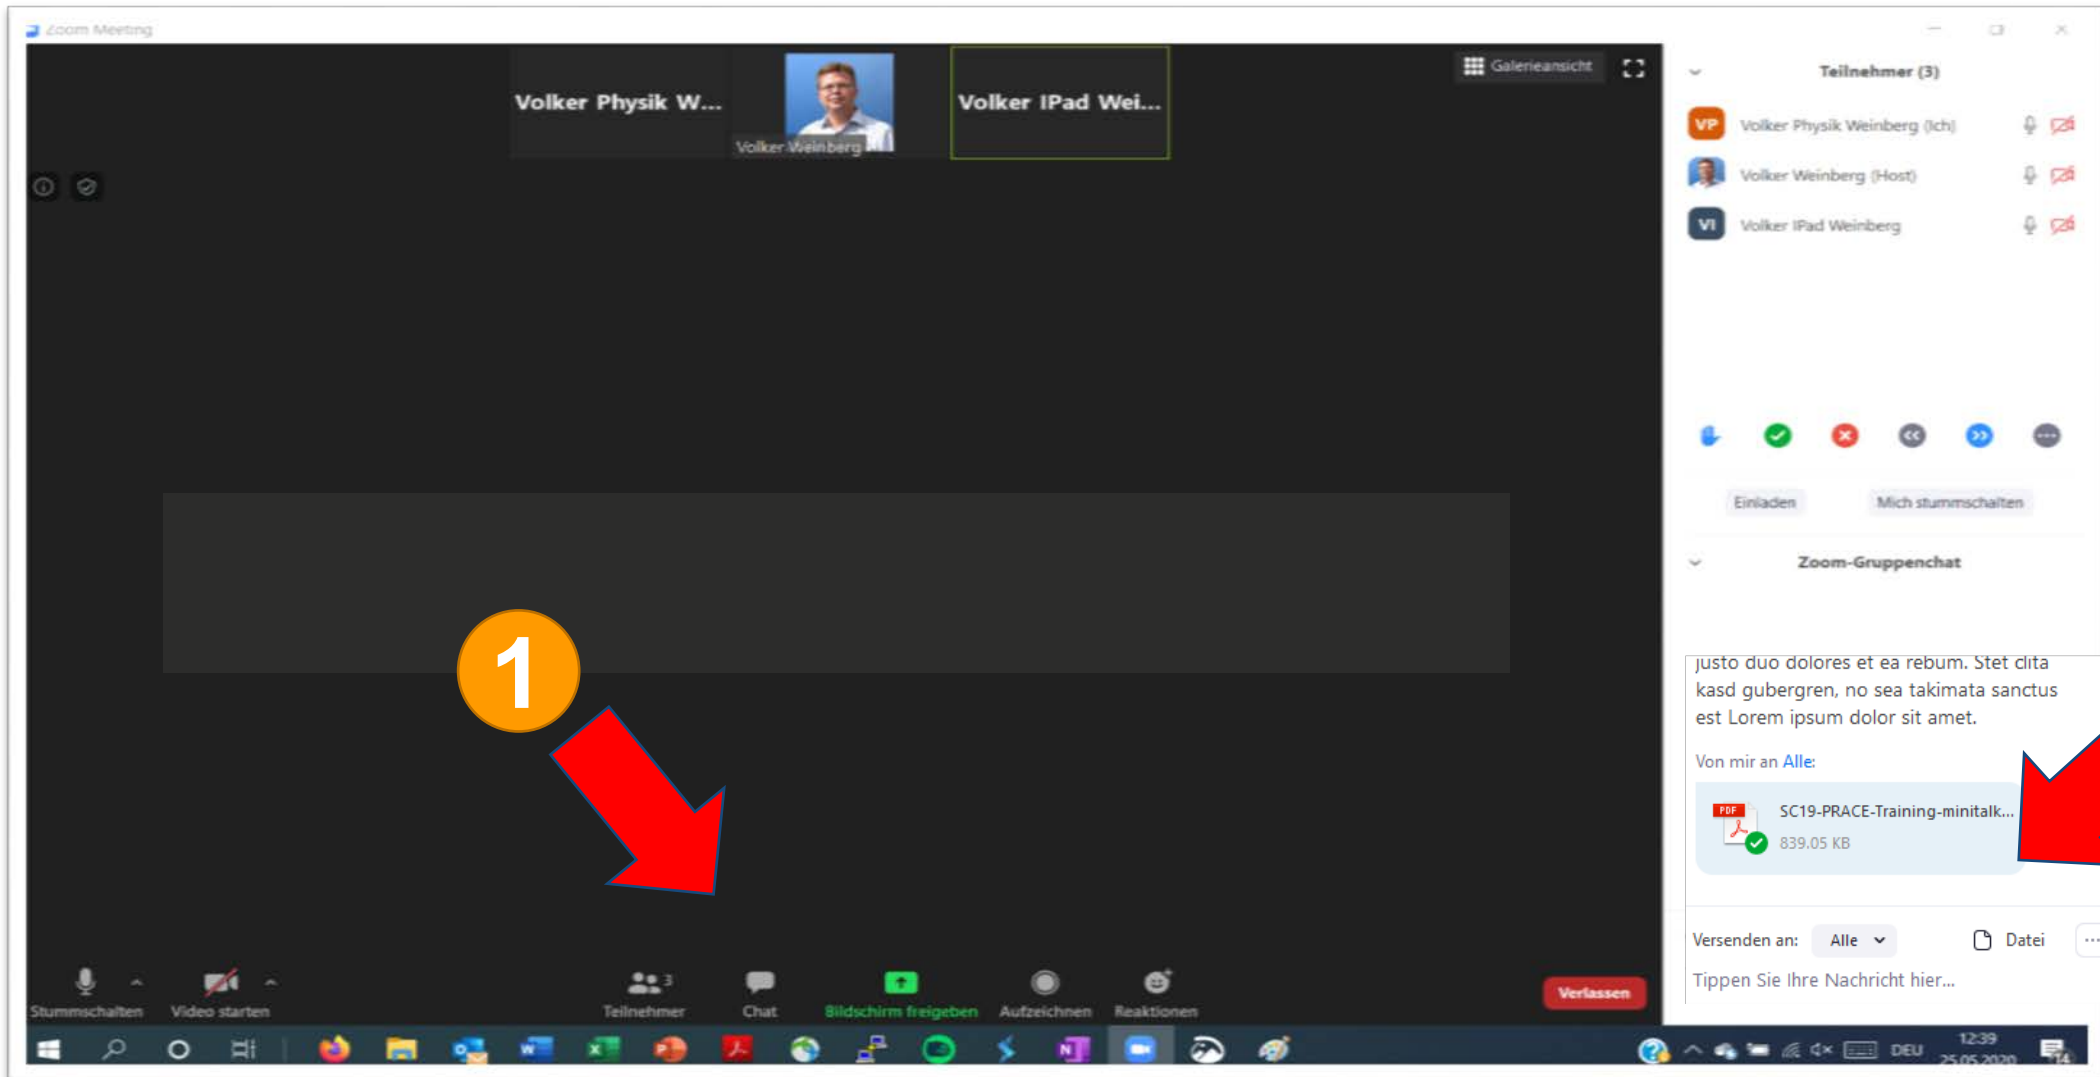


- If you do not mind, please show your video when asking questions.
- **Push to Talk:**  
The Push to Talk feature allows you to remain muted throughout your Zoom meeting and hold down the spacebar when you want to be unmuted and talk.
- **Instant Feedback:**



# zoom

## Chat and Share Files



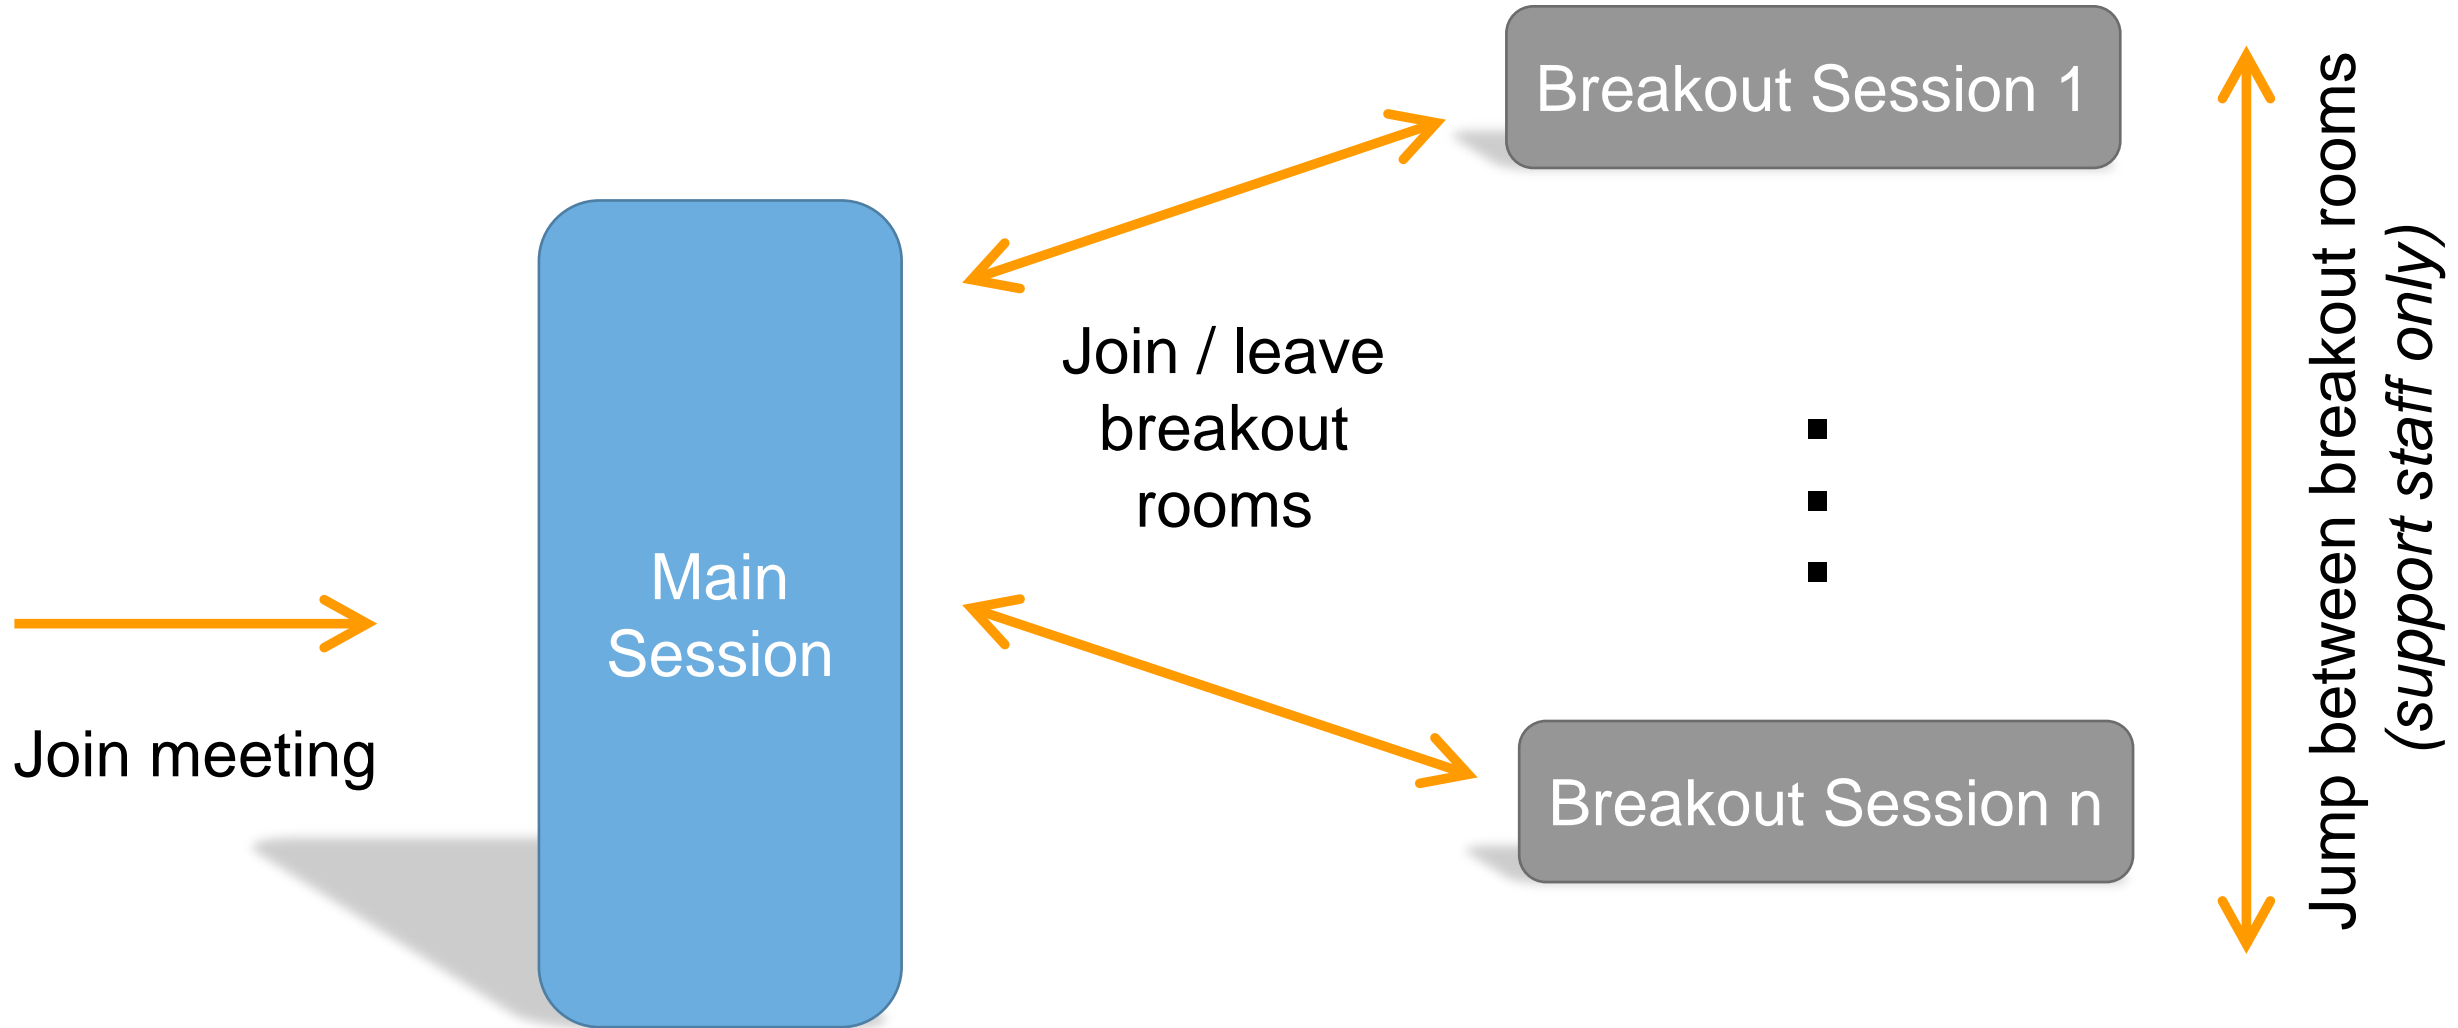
- Breakout rooms allow you to split your Zoom meeting in up to 50 separate sessions. The meeting host can choose to split the participants of the meeting into these separate sessions automatically or manually, and can switch between sessions at any time.
- We will have 2 Breakout Rooms:
  - Breakout Session 1
  - Breakout Session 2
- Participants will be automatically assigned to one of these breakout rooms during the exercise sessions.



# Breakout Rooms


	Main Session	Breakout Room 1	Breakout Room 2
<b>1st day: Fundamentals of Deep Learning for Computer Vision</b>	Juan	Juan	Volker
<b>2nd day: Fundamentals of Deep Learning for Multiple Data Types</b>	Juan	Juan	Momme
<b>3rd day: Fundamentals of Accelerated Computing with OpenACC</b>	Volker	Volker	Momme



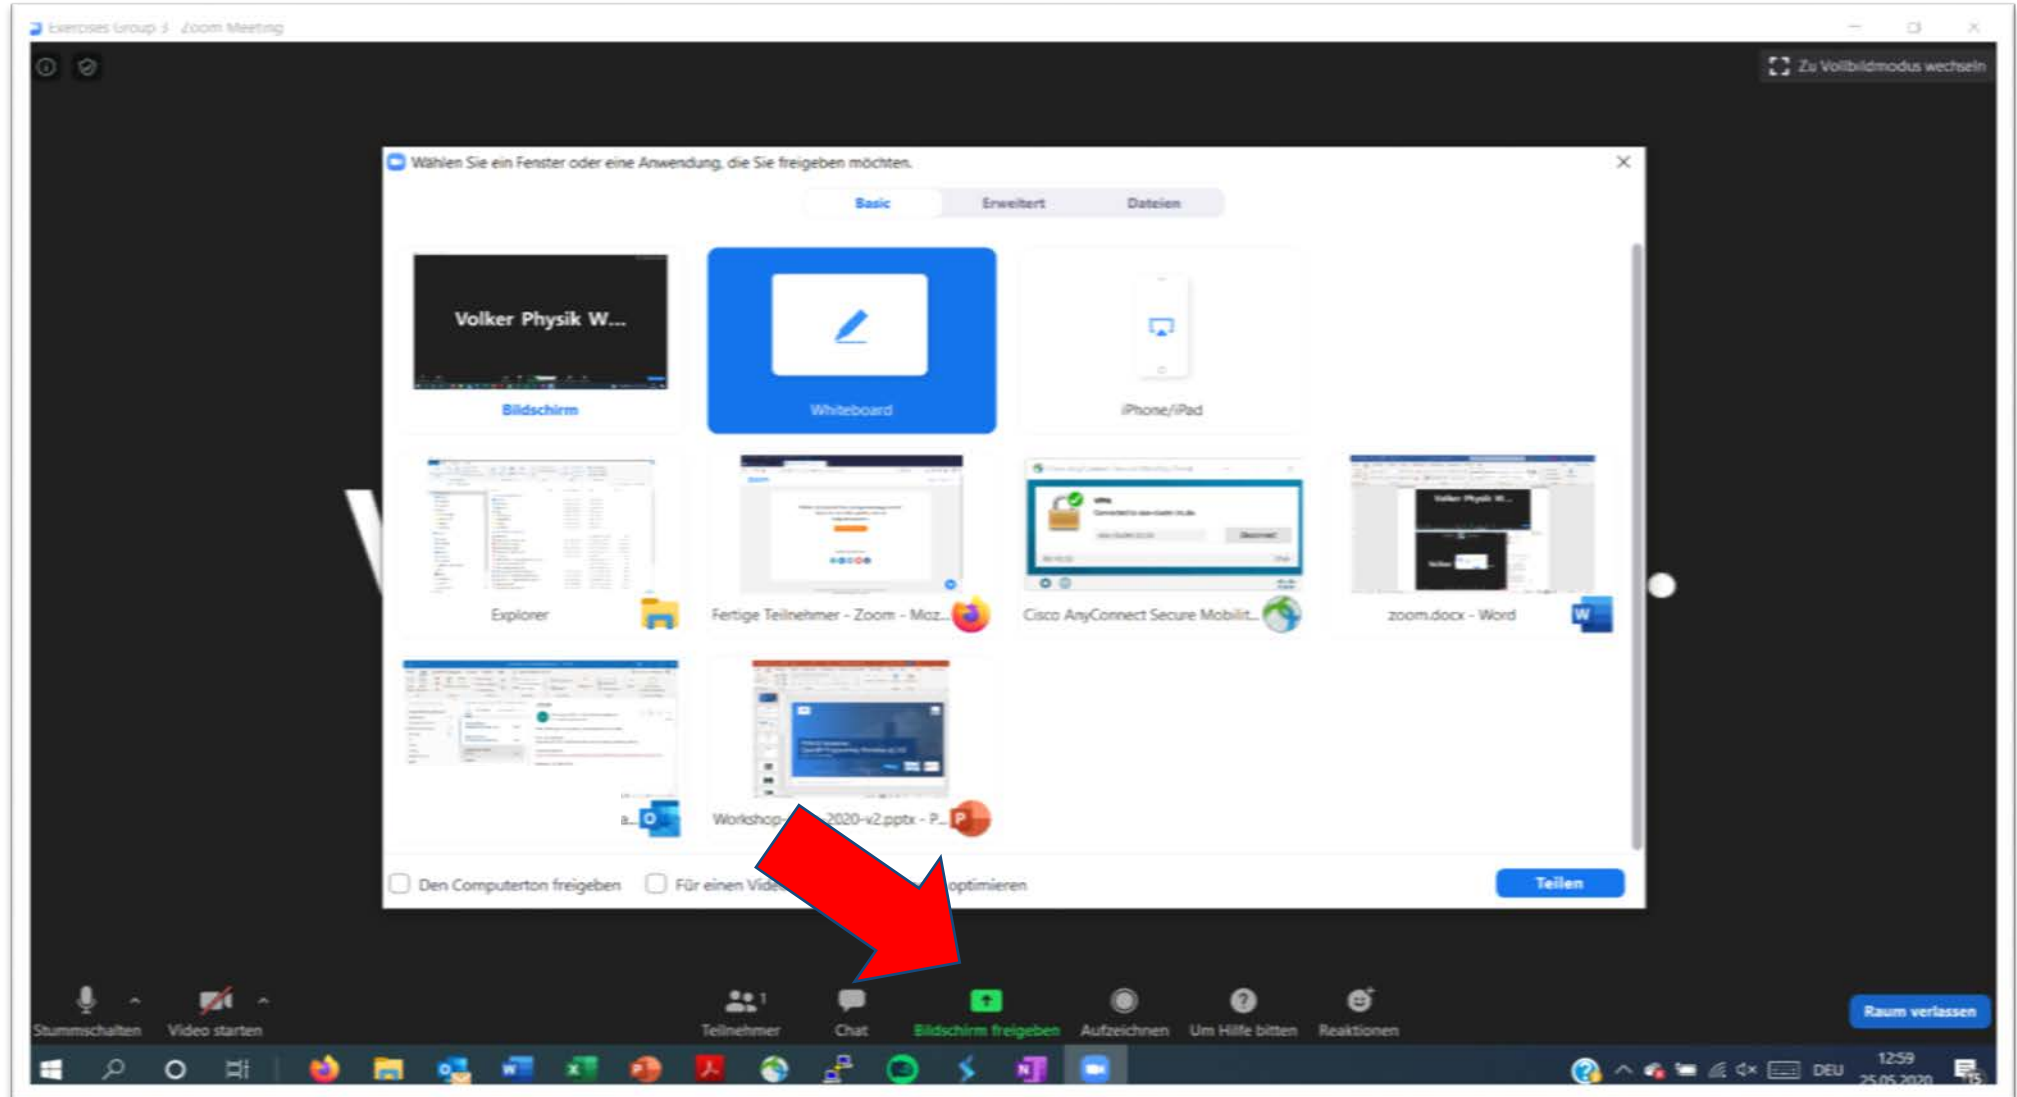


# Breakout Rooms: Ask for help as a participant...

... if you are in a breakout room and your personal tutor is not available



... participants  
can share  
screens /  
whiteboard in  
breakout  
rooms per  
default



And now ...



DEEP  
LEARNING  
INSTITUTE



**Enjoy the workshop!**