



DEEP
LEARNING
INSTITUTE

PRACE Workshop: Deep Learning and GPU programming workshop

7 – 10 September 2020



VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

Overview



- The workshop is co-organized by LRZ, CSC, IT4Innovations and NVIDIA Deep Learning Institute (DLI) for the Partnership for Advanced Computing in Europe (PRACE).
- NVIDIA Deep Learning Institute (DLI) offers hands-on training for developers, data scientists, and researchers looking to solve challenging problems with deep learning.
- This 4-days workshop offered for the first time online combines lectures about fundamentals of Deep Learning for Multiple Data Types and Multi-GPUs with lectures about Accelerated Computing with OpenACC and CUDA C/C++.
- 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, how to effectively parallelize training of deep neural networks on Multi-GPUs and how to accelerate your applications with OpenACC and CUDA C/C++.
- 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.



PRACE Training



LRZ as part of the Gauss Centre for Supercomputing (GCS), CSC and IT4Innovations belong to the 14 **PRACE Training Centres** that started in 2012-2017-2020:

- Barcelona Supercomputing Center (Spain)
- CINECA Consorzio Interuniversitario (Italy)
- CSC – IT Center for Science Ltd (Finland)
- EPCC at the University of Edinburgh (UK)
- Gauss Centre for Supercomputing (Germany)
- Maison de la Simulation (France)
- GRNET – Greek Research and Technology Network (Greece)
- ICHEC – Irish Centre for High-End Computing (Ireland)
- IT4I – National Supercomputing Center VSB Technical University of Ostrava (Czech Republic)
- SURFsara (The Netherlands)
- TU Wien – VSC Research Center (Austria)
- University ANTWERPEN – VSC & CÉCI (Belgium)
- University of Ljubljana – HPC Center Slovenia (Slovenia)
- Swedish National Infrastructure for Computing (SNIC) (Sweden)



Univerza v Ljubljani



Irish Centre for High-End Computing



IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER



Mission: Serve as **European hubs and key drivers of advanced high-quality training** for researchers working in the computational sciences.

<http://www.training.prace-ri.eu/>



IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER



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.



VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

NVIDIA DEEP LEARNING INSTITUTE

Learn more at
www.nvidia.co.uk/dli

ONLINE SELF-PACED TRAINING

Deep Learning, Accelerated Computing and Accelerated Data Science courses online

TEACHING KITS

- Qualified university educators can download courseware across deep learning, accelerated computing, and robotics.
- Kits include lecture materials, GPU cloud resources, free access to self-paced DLI courses, and more.

<https://developer.nvidia.com/teaching-kits>

UNIVERSITY AMBASSADOR PROGRAM

- Qualified faculty and researchers can get certified to teach DLI workshops to their students at no cost.

GTC 5.-9. OCTOBER 2020

- Free for academia. DLI workshops cost 99 USD.
- www.nvidia.com/gtc
- Use reference code NYMDIERINGER, it provides industry customers 10 USD discount on the conference passes (not applicable to alumni rates or DLIs)

Lecturers



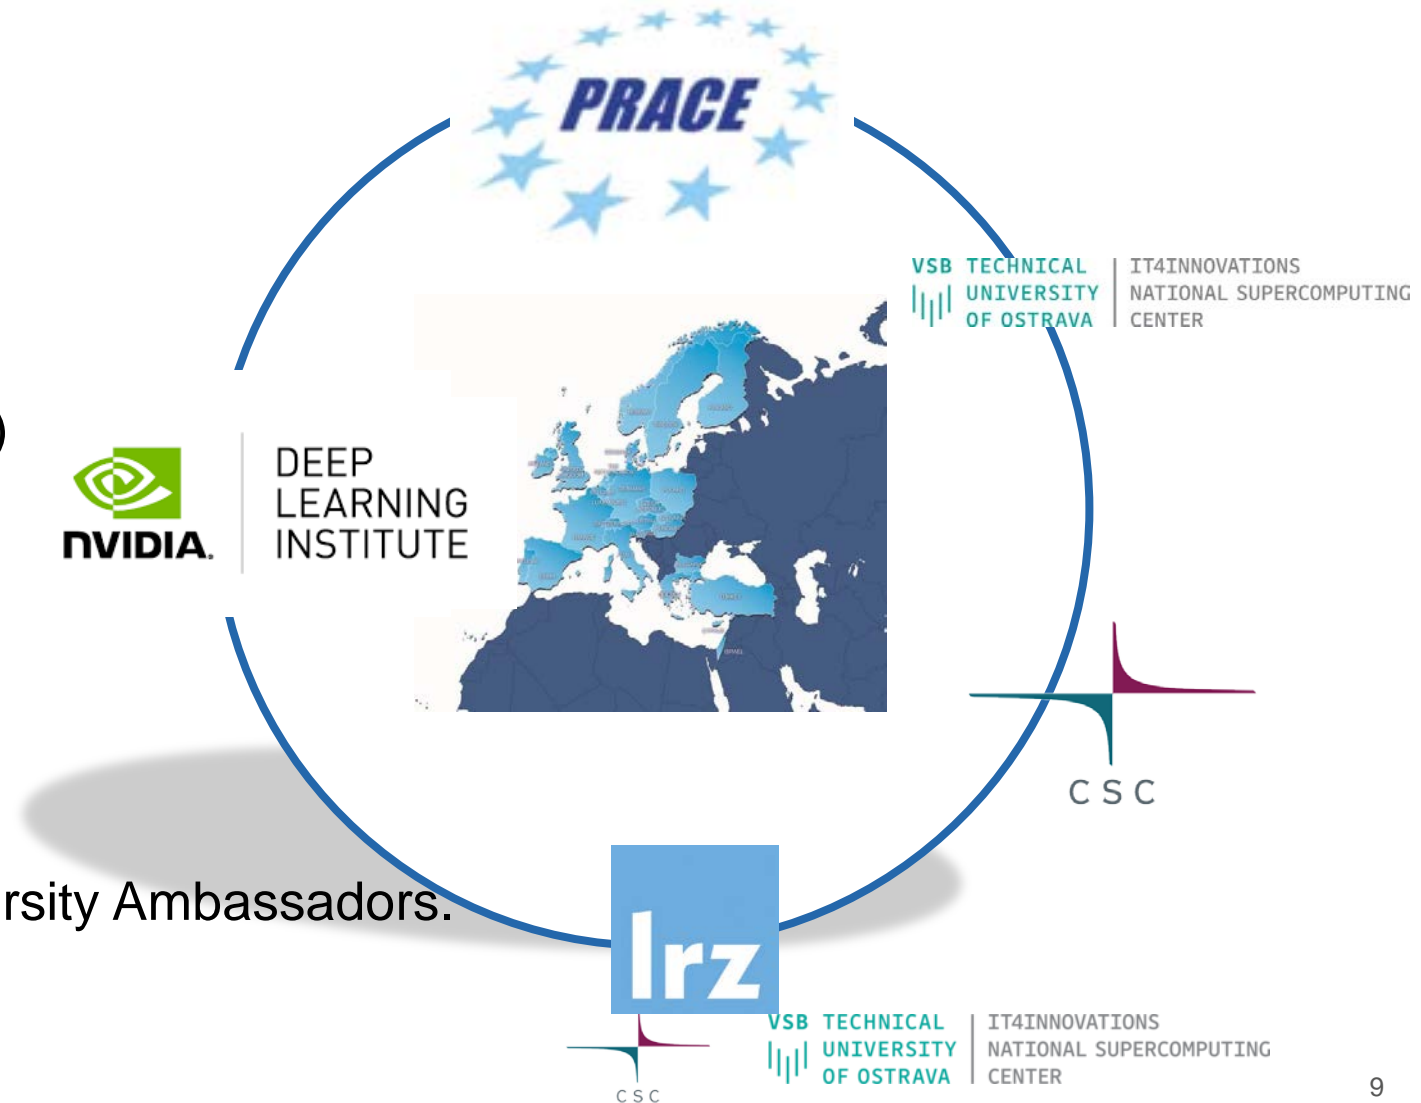
- **Lecturers:**

- Dr. Momme Allalen (LRZ)
- PD Dr. Juan Durillo Barrionuevo (LRZ)
- Dr. Volker Weinberg (LRZ)
- Georg Zitzlsberger (IT4Innovations)

All instructors are NVIDIA certified University Ambassadors.



DEEP
LEARNING
INSTITUTE



Lecturers



DEEP
LEARNING
INSTITUTE



	Lecturer	Assistant
1st day: Fundamentals of Deep Learning for Multiple Data Types	Juan	Georg
2nd day: Fundamentals of Accelerated Computing with OpenACC	Volker	Momme
3rd day: Fundamentals of Accelerated Computing with CUDA C/C++	Momme	Volker
4th day: Fundamentals of Deep Learning for Multi-GPUs	Georg	Juan



VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

1st day: Fundamentals of Deep Learning for Multiple Data Types



- The 1st 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.

2nd day: Fundamentals of Accelerated Computing with OpenACC



- On the 2nd 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.

3rd day: Fundamentals of Accelerated Computing with CUDA C/C++



- The CUDA computing platform enables the acceleration of CPU-only applications to run on the world's fastest massively parallel GPUs. On the 3rd day you experience C/C++ application acceleration by:
 - Accelerating CPU-only applications to run their latent parallelism on GPUs
 - Utilizing essential CUDA memory management techniques to optimize accelerated applications
 - Exposing accelerated application potential for concurrency and exploiting it with CUDA streams
 - Leveraging command line and visual profiling to guide and check your work
- Upon completion, you'll be able to accelerate and optimize existing C/C++ CPU-only applications using the most essential CUDA tools and techniques. You'll understand an iterative style of CUDA development that will allow you to ship accelerated applications fast.

4th 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 4th 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 Deep Learning for Multiple Data Types



- 10:00-10:20 Welcome and Intro
- 10:20-12:00 Introduction to CNNs and Object Segmentation
- 12:00-13:00 Lunch Break**
- 13:00-14:20 Word Generation with RNNs
- 14:20-14:30 Coffee Break**
- 14:30 Group Picture**
- 14:30-15:45 Image Captioning by Combining RNNs and CNNs
- 15:45-16:00 Q&A



**All times are in
EEST=CEST+1!**

Tentative Agenda Day 2: **Fundamentals of Accelerated Computing with OpenACC**



10:00-10:15 Intro and 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
EEST=CEST+1!**

Tentative Agenda Day 3: Fundamentals of Accelerated Computing with CUDA C/C++



DEEP
LEARNING
INSTITUTE



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

12:00-13:00 Lunch Break

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



**All times are in
EEST=CEST+1!**

Tentative Agenda Day 4: Fundamentals of Deep Learning for Multi-GPUs



10:00-10:15 Intro

10:15-12:00 Stochastic Gradient Descent


12:00-13:00 Lunch Break

13:00-14:20 Introduction to Distributed Training

14:20-14:30 Coffee Break

14:30-15:45 Algorithmic Challenges of Distributed SGD

15:45-16:00 Q&A



**All times are in
EEST=CEST+1!**

Workshop Webpage



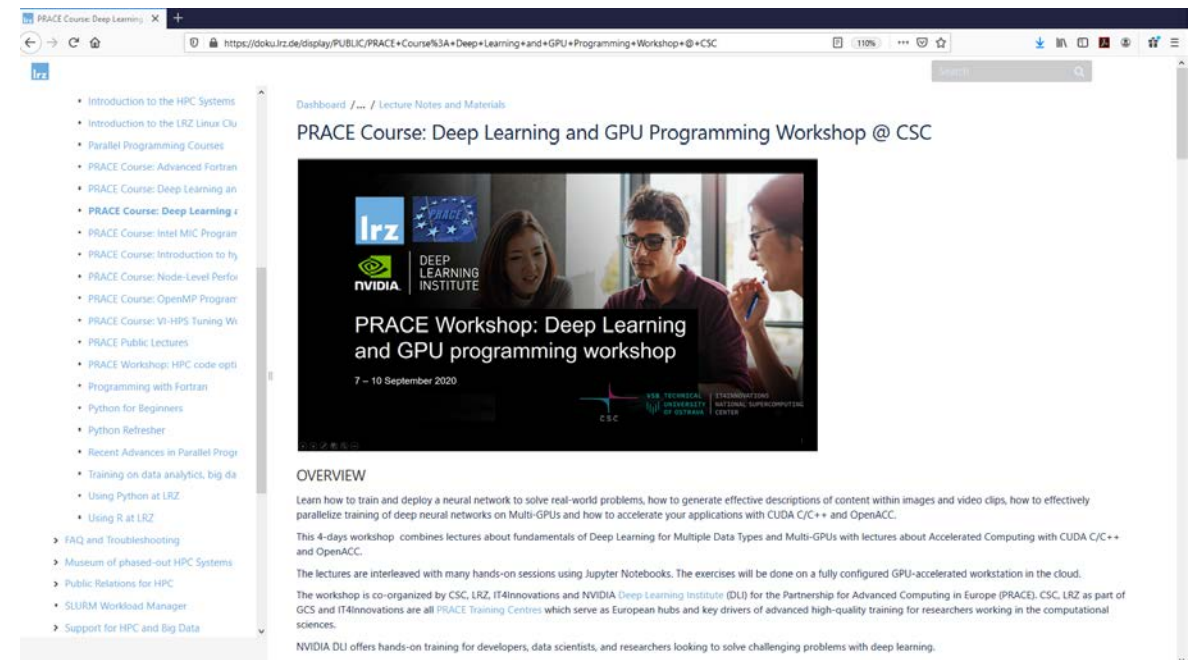
- All slides will be made available during the workshop under:

- <https://tinyurl.com/dl-gpu-workshop-csc>



- 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/81313420308?pwd=MGllDVEOHZ0cFFMd29ZWldxeVJEZz09>

Meeting ID: 813 1342 0308


Passcode: 9wDTcu

If you have any audio problems, you can also join via phone:

Meeting ID: 813 1342 0308

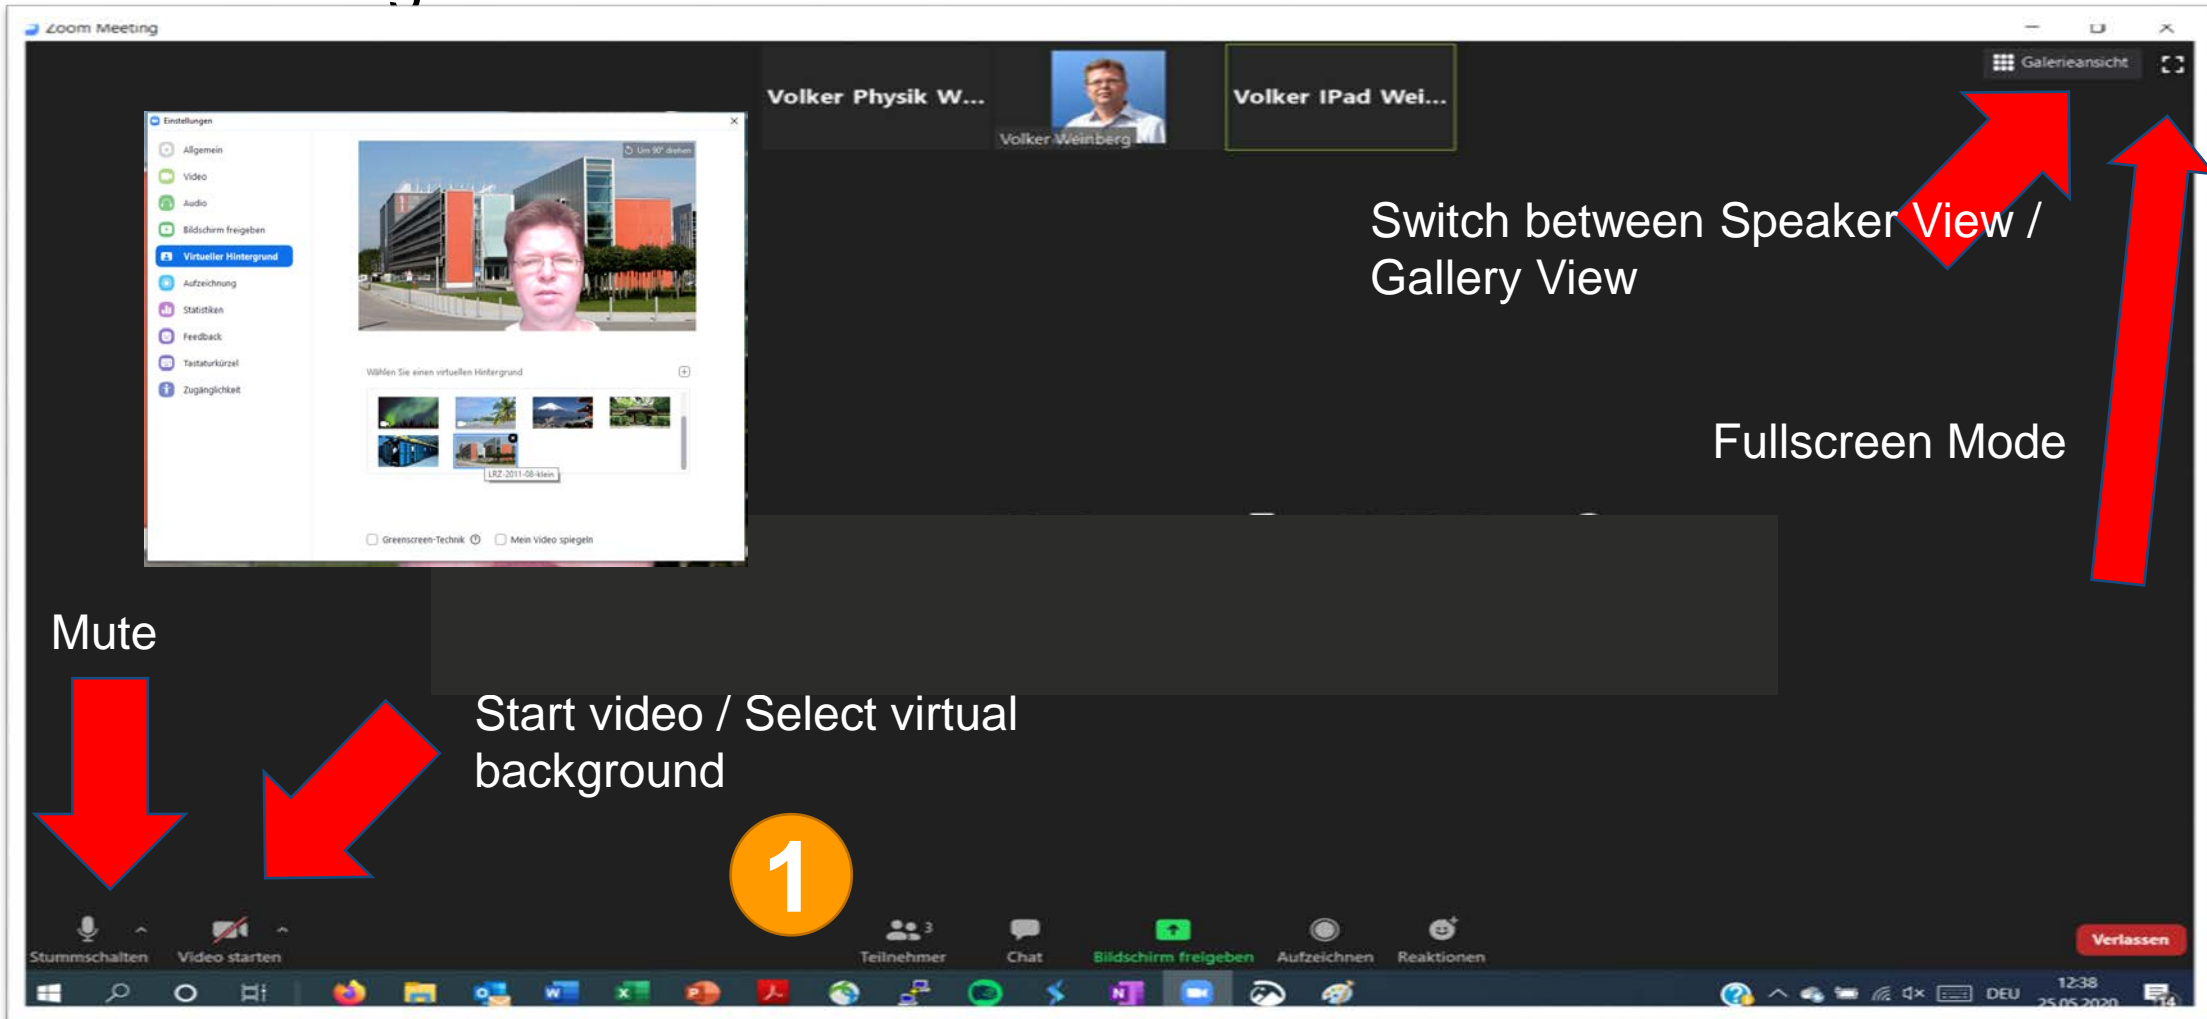
Passcode: 428546

Find your local number: <https://us02web.zoom.us/u/keIFzHf9GO>



- 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. 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 three participants: Volker Physik Weinberg, Volker Weinberg (Host), and Volker iPad Weinberg. The bottom toolbar includes buttons for 'Stummschalten', 'Video starten', 'Teilnehmer', 'Chat', 'Bildschirm freigeben', 'Aufzeichnen', 'Reaktionen', and 'Verlassen'. The 'Teilnehmer' button is highlighted with a red arrow and the number 1. The 'Reaktionen' button is also highlighted with a red arrow and the number 2. A tooltip for 'Weitere nonverbale Rückmeldungen' is visible, showing various reaction icons like thumbs up, thumbs down, clapping, and a clock.



- 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



The screenshot shows a Zoom meeting window with a chat panel on the right. The chat panel displays a message with a PDF file attachment. A red arrow labeled '1' points to the chat area, and another red arrow labeled '2' points to the file attachment. The file name is 'SC19-PRACE-Training-minitalk...' and the size is 839.05 KB. The chat area also shows a text message: 'justo duo dolores et ea rebum. Stet dita kasd gubergren, no sea takimata sanctus est Lorem ipsum dolor sit amet.' and a 'Versenden an: Alle' dropdown menu.

And now ...



DEEP
LEARNING
INSTITUTE



Enjoy the workshop!



VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER