

A white wolf is walking towards the viewer in a snowy, futuristic landscape. The background is dark blue with vertical light beams and digital particles, creating a high-tech, AI-themed environment.

LUMI

First AI training
job on LUMI

Mats Sjöberg, Oskar Taubert – CSC – IT Center for Science, Finland

Accessing LUMI via SSH



- No passwords
→ SSH key pair **required**
- register your public key in your MyAccessID user profile
 - Finnish users with MyCSC linked account should use MyCSC for SSH keys!

```
ssh -i <path-to-private-key> \  
    <username>@lumi.csc.fi
```

<https://docs.lumi-supercomputer.eu/firststeps/>

A terminal window showing an SSH session. The prompt is 'skata:~ \$' and the command is 'ssh mvsjobber@lumi.csc.fi'. The terminal displays a large ASCII art logo of the word 'LUMI' in a stylized, blocky font. Below the logo, it says 'The Supercomputer of the North'. At the bottom, there is a section titled 'User guide and support' with two URLs: 'https://docs.lumi-supercomputer.eu' and 'https://lumi-supercomputer.eu/user-support'. The prompt at the bottom is 'uan01:~ \$'.

Pinned Apps



Home Directory



Compute node shell



Login node shell



Desktop



Cloud storage
configuration



Disk quotas



Project view



Active Jobs



Jupyter



Jupyter for courses



Julia-Jupyter



MATLAB



MLflow



TensorBoard



Visual Studio Code

Setting up the Software Environment

LUMI

<https://pytorch.org/>

PyTorch Build	Stable (2.5.1)			Preview (Nightly)		
Your OS	Linux		Mac		Windows	
Package	Conda		Pip		LibTorch	Source
Language	Python			C++ / Java		
Compute Platform	CUDA 11.8	CUDA 12.1	CUDA 12.4	ROCm 6.2		CPU
Run this Command:	pip3 install torch torchvision torchaudio --index-url https://download.pytorch.org/whl/rocm6.2					

```
! environment-minimal.yml
1  channels:
2    - pytorch
3    - conda-forge
4    - defaults
5  dependencies:
6    - accelerate=0.29.1
7    - datasets=2.18.0
8    - python=3.10.14
9    - pytorch=2.2.2
10   - transformers=4.39.3
11
```

~47k files



**bad for the
shared network
filesystem**

Setting up the Software Environment

- Singularity/Apptainer containers recommended for ML applications on LUMI
- Similar to Docker, better suited for multi-user supercomputers
 - Isolated environment
 - Easier to manage complex software dependencies
- For now we'll provide a pre-installed container
 - Later lectures will discuss how to create your own containers

python



singularity exec \$CONTAINER python

Setting up the Software Environment



Set up the software environment



```
$ module purge
$ module use /appl/local/containers/ai-modules
$ module load singularity-AI-bindings
```

Launch the PyTorch container

```
$ CONTAINER=/appl/local/containers/sif-images/lumi-pytorch-rocm-6.2.4-
python-3.12-pytorch-v2.6.0.sif
$ singularity exec $CONTAINER python
```


Don't run computation on the login node
– use the Slurm job scheduler!

LUMI



SLURM is used to reserve resources and submit scripts for running on the compute nodes

Batch system



Training on a single Graphics Compute Die

LUMI

SLURM batch script (run.sh)

```
#!/bin/bash
```

```
#SBATCH --account=project_XXXXXXXXX
```

```
#SBATCH --partition=small-g
```

```
#SBATCH --gpus-per-node=1
```

```
#SBATCH --ntasks-per-node=1
```

```
#SBATCH --cpus-per-task=7
```

```
#SBATCH --mem-per-gpu=60G
```

```
#SBATCH --time=1:00:00
```

What resources requested?

- 1 GPU = 1/8 of node
- Use also $\leq 1/8$ of CPU cores and memory

```
module purge
```

```
module use /appl/local/containers/ai-modules
```

```
module load singularity-AI-bindings
```

```
CONTAINER=/path/to/pytorch.sif
```

What software to load?

```
srun singularity exec $CONTAINER python my_pytorch_script.py
```

Available GPU partitions

- **standard-g**
≤ 48h, whole nodes only, max 1024 nodes/job
- **small-g**
≤ 72h, individual GCDs, max 4 nodes/job
- **dev-g**
≤ 3h, individual GCDs, max 32 nodes/job, *max 2 jobs running*

Useful SLURM commands

Submit the SLURM batch script

```
$ sbatch run.sh  
Submitted batch job 987654
```

Check the SLURM queue

```
$ squeue --me  
  
JOBID    PARTITION NAME    ST TIME    NODES NODELIST  
987654   small-g   run.sh R   0:18      1 compute_node
```

Cancel a job

```
$ scancel 987654
```

Useful SLURM commands



Read job outputs

```
$ tail -f slurm-987654.out
```

Access compute node for diagnostics

```
$ srun --overlap --pty --jobid=987654 bash  
[compute_node] $ rocm-smi
```

Useful SLURM commands for monitoring job progress

LUMI

Check the status of partitions

```
$ sinfo -s
```

PARTITION	AVAIL	TIMELIMIT	NODES(A/I/O/T)	...
debug	up	30:00	8/0/0/8	nid[002595-002597]
interactive	up	8:00:00	4/0/0/4	nid[002502,002507,002594,002599]
q_fiqci	inact	15:00	0/1/0/1	nid002598
q_industry	up	15:00	0/1/0/1	nid002598
q_nordiq	up	15:00	0/1/0/1	nid002503
small	up	3-00:00:00	280/6/20/306	nid[002280-002499,002508-002593]
standard	up	2-00:00:00	1493/111/124/172	nid[001000-002279,002600-003047]
dev-g	up	3:00:00	30/17/1/48	nid[005002-005025,007954-007977]
small-g	up	3-00:00:00	193/2/3/198	nid[005026-005123,007852-007951]
standard-g	up	2-00:00:00	2555/23/150/2728	nid[005124-007851]
largemem	up	1-00:00:00	1/5/0/6	nid[000101-000106]
lumid	up	4:00:00	1/6/1/8	nid[000016-000023]

A = allocated

I = idle

O = other states

T = total

Monitoring Progress with TensorBoard

LUMI



- HuggingFace trainer automatically writes logs for TensorBoard
- You can start the UI in the LUMI web interface
- For the exercises use path:
`/scratch/project_465002178/USERNAME/runs/`

https://pytorch.org/tutorials/recipes/recipes/tensorboard_with_pytorch.html

