LUMI AI Guide

AI group – LUMI User Support Team

State of AI use on LUMI

LUMI

- LUMI end-user survey 2024:
  - 29% of LUMI users state AI as main field of computation
  - 41% use Machine Learning frameworks on LUMI in one way or another
  - PyTorch by far most popular framework, followed by Scikit-learn, TensorFlow, Jax, ...
  - Number of AI users increases every year

- LUMI was not designed as a supercomputer optimized for AI tasks
  - … but we are working on making it more and more user friendly

Challenges for new AI users on LUMI

LUMI

- Software installations (EasyBuild modules, singularity containers, virtual environments, ...)
- AMD GPUs instead of more commonly used NVidia hardware
- "Many small files problem" - compatibility with Lustre file system
- Scaling training jobs to multiple GPUs and nodes
- Monitoring and profiling training jobs
- ...

Upcoming Events

26.3.2025
Online
LUMI user coffee breaks

27.-28.3.2025
CSC, Espoo, Finland
GPU Programming with HIP

2.4.2025
Pikku-Finlandia, Helsinki, Finland and online
LUMI AI Factory Launch

12.-16.5.2025
Oslo, Norway
Hackathon: Optimizing for

Moving your AI training jobs to LUMI: A Hands-On Workshop

4.-5.2.2025
Espoo, Finland

Join our two-day workshop, "Getting Started with AI on LUMI," designed to familiarize you with the capabilities of the LUMI supercomputer for artificial intelligence applications. This workshop is ideal for those looking to transition from smaller-scale computing environments like laptops, workstations, or cloud VMs to the robust, GPU-intensive LUMI platform.

Participants are invited to **bring their own AI training scripts to the workshop**, where they will receive personalized support to adapt and run them on LUMI's advanced GPU system. Whether you aim to leverage a single GPU or scale up to multiple GPUs, our workshop will provide valuable insights and practical skills to enhance your AI projects with LUMI's powerful computing infrastructure.

Event Info

Dates: 4.-5. February 2025
Time: 9:00 – 17 EET each day
Location: Espoo, Finland
Organizer: LUMI User Support Team (LUST), EuroCC National Competence Centers (NCCs) in Finland and GreenNLP

LUMI **Documentation**

Search

Home | First steps | Hardware | Run jobs | Software | Developing | Storage | Tutorials | Help desk

Welcome

Welcome to the LUMI supercomputer user guide. To navigate this guide, select a category from the navigation bar at the top of the page or use the search function.

You have not connected to LUMI yet? Please visit the first steps section to get started.

→ First steps

LUMI helpdesk | LUMI status | LUMI events | LUMI training materials

Discover the LUMI Hardware

→ GPU partition
→ CPU partition
→ Visualization partition

Submitting a Job

→ Available Slurm partitions
→ Example GPU jobs
→ Example CPU jobs

Storage

→ Data storage options
→ Using Lustre efficiently
→ Object storage

Overview

Training Material

The LUMI training material is hosted at lumi-supercomputer.github.io/LUMI-training-materials.

It contains all slides and recordings of all past LUMI training events, including introductory and advanced courses, AI workshops, hackathons, and profiling courses. The training material also provides recordings of the LUMI user coffee break talks and user update events.

LUMI AI Guide

The LUMI AI Guide is hosted at github.com/Lumi-supercomputer/LUMI-AI-Guide.

This guide is designed to assist users in migrating their machine learning applications from smaller-scale computing environments to LUMI. We will walk you through a detailed example of training an image classification model using PyTorch's Vision

LUMI-AI-Guide

Public

Edit Pins | Unwatch 7 | Fork 7 | Starred 28

main | Go to file t | Code

gregordecristoforo use correct link to python script 1254073 · 5 days ago 224 Commits

1-quickstart	add dots to table of content	2 weeks ago
2-setting-up-environment	add dots to table of content	2 weeks ago
3-file-formats	add dots to table of content	2 weeks ago
4-data-storage	add dots to table of content	2 weeks ago
5-multi-gpu-and-node	add dots to table of content	2 weeks ago
6-monitoring-and-profiling	add dots to table of content	2 weeks ago
7-TensorBoard-visualization	add dots to table of content	2 weeks ago
8-MLflow-visualization	add dots to table of content	2 weeks ago
9-Wandb-visualization	add dots to table of content	2 weeks ago
assets/images	Wandb visualization	last month

About

The LUMI AI Guide is designed to assist users in migrating their machine learning applications from smaller-scale computing environments to the LUMI supercomputer.

ai deep-learning-tutorial lumi

Readme

Unknown, Unknown licenses found

Cite this repository

Activity

Custom properties

28 stars

7 watching

7 forks

Report repository

LUMI AI Guide

LUMI

- Designed to assist users in migrating their machine learning applications from smaller-scale computing environments to LUMI

- Includes an A-Z example of training an image classification model using a Vision Transformer (VIT) on the ImageNet dataset on LUMI

Table of contents

The guide is structured into the following sections:

Requirements

LUMI

- A basic understanding of machine learning concepts and Python programming. The guide will focus primarily on aspects specific to training models on LUMI.
- An active user account on LUMI and familiarity with its basic operations.
- If you wish to run the included examples, you need to be part of a project with GPU hours on LUMI.

Some highlights

LUMI

Chapter 2. Setting up your own environment

- Provided containers on LUMI
- Interacting with a containerized environment
- Installing additional Python packages in a container
- Using Custom container images

•Chapter 3. File formats for training data

- Squashfs
- HDF5
- LMDB
- Performance benchmarks
- Scripts to convert data from one format to another

Some highlights

LUMI

Chapter 4. Data Storage Options

- LUMI-P
- LUMI-F
- RAMfs
- File striping

Chapter 5. Multi-GPU and Multi-Node Training

- PyTorch DDP
- Torchrun
- DeepSpeed
- Optimal CPU-GPU bindings
- RCCL environment variables

Chapter 6. Monitoring and Profiling jobs

- Monitoring jobs with rocm-smi
- PyTorch profiler
- AMD profilers

```
decristoforo@nid005216:/project/project_462000002/LUMI-AI-example$ rocm-smi


======================================== ROCm System Management Interface ========================================
================================================== Concise Info ==================================================
Device  [Model : Revision]    Temp    Power   Partitions      SCLK     MCLK     Fan  Perf    PwrCap  VRAM%  GPU%
        Name (20 chars)       (Edge)  (Avg)   (Mem, Compute)
==================================================================================================================
0       [0x0b0c : 0x00]       60.0°C  475.0W  N/A, N/A        1445Mhz  1600Mhz  0%   manual  500.0W  12%    100%
        AMD INSTINCT MI200 (
1       [0x0b0c : 0x00]       64.0°C  N/A     N/A, N/A        1400Mhz  1600Mhz  0%   manual  0.0W    12%    100%
        AMD INSTINCT MI200 (
2       [0x0b0c : 0x00]       55.0°C  476.0W  N/A, N/A        1440Mhz  1600Mhz  0%   manual  500.0W  12%    100%
        AMD INSTINCT MI200 (
                                              N/A, N/A        1390Mhz  1600Mhz  0%   manual  0.0W    12%    100%
                                              N/A, N/A        1430Mhz  1600Mhz  0%   manual  500.0W  12%    100%
                                              N/A, N/A        1385Mhz  1600Mhz  0%   manual  0.0W    12%    100%
                                              N/A, N/A        1460Mhz  1600Mhz  0%   manual  500.0W  12%    100%
                                              N/A, N/A        1465Mhz  1600Mhz  0%   manual  0.0W    12%    100%
==================================================================================================================
= End of ROCm SMI Log ============================================================================================
```

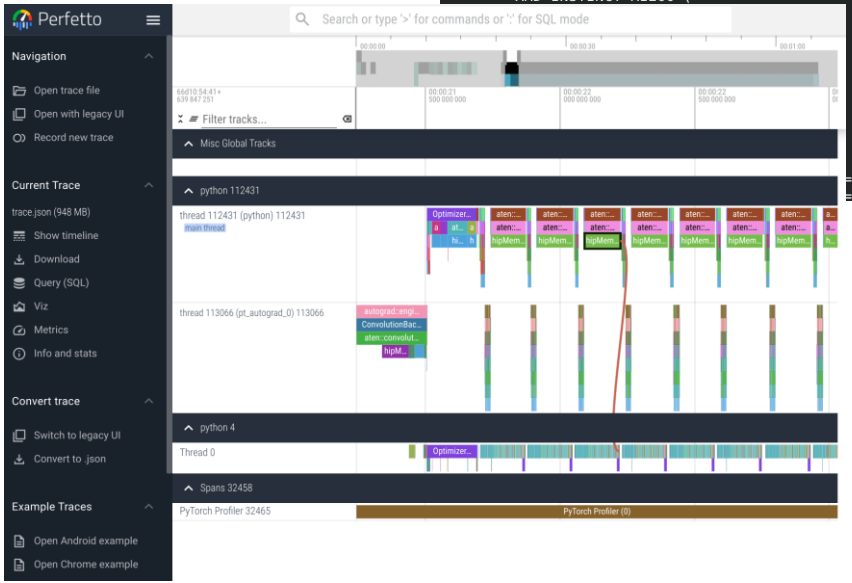

Chapter 7 – 9: Visualization tools: TensorBoard, MLflow, Wandb

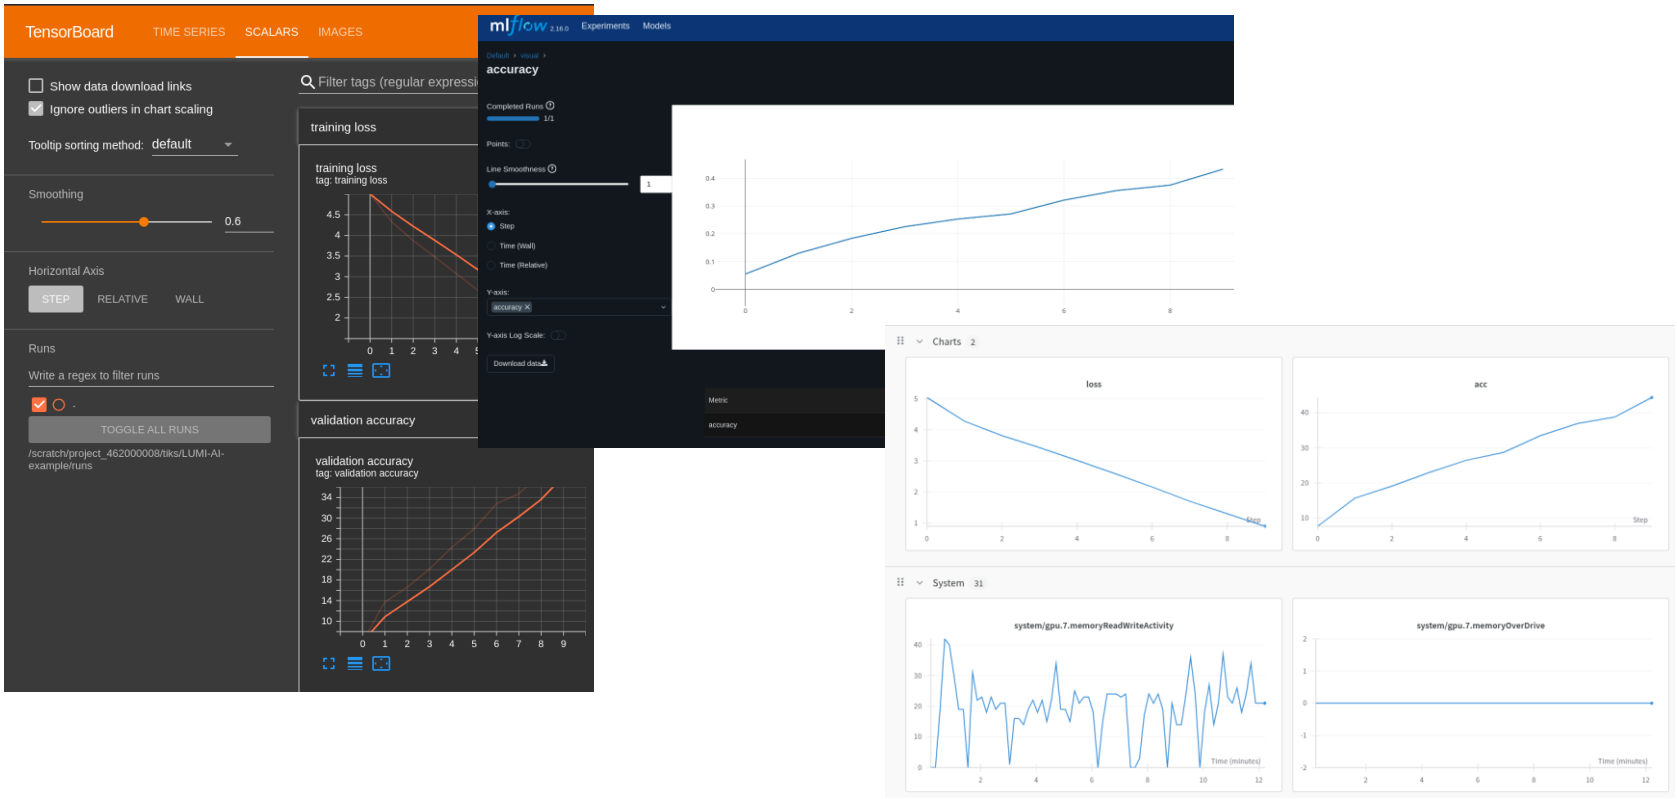

Code examples for all chapters

You can run included examples yourself!

- set_up_environment.sh script copies container and training data to your directory
- Python and Slurm scripts included for every chapter
- ViT (Vision Transformer) example starts in first chapter on a single GPU
- Scripts are modified for every chapter

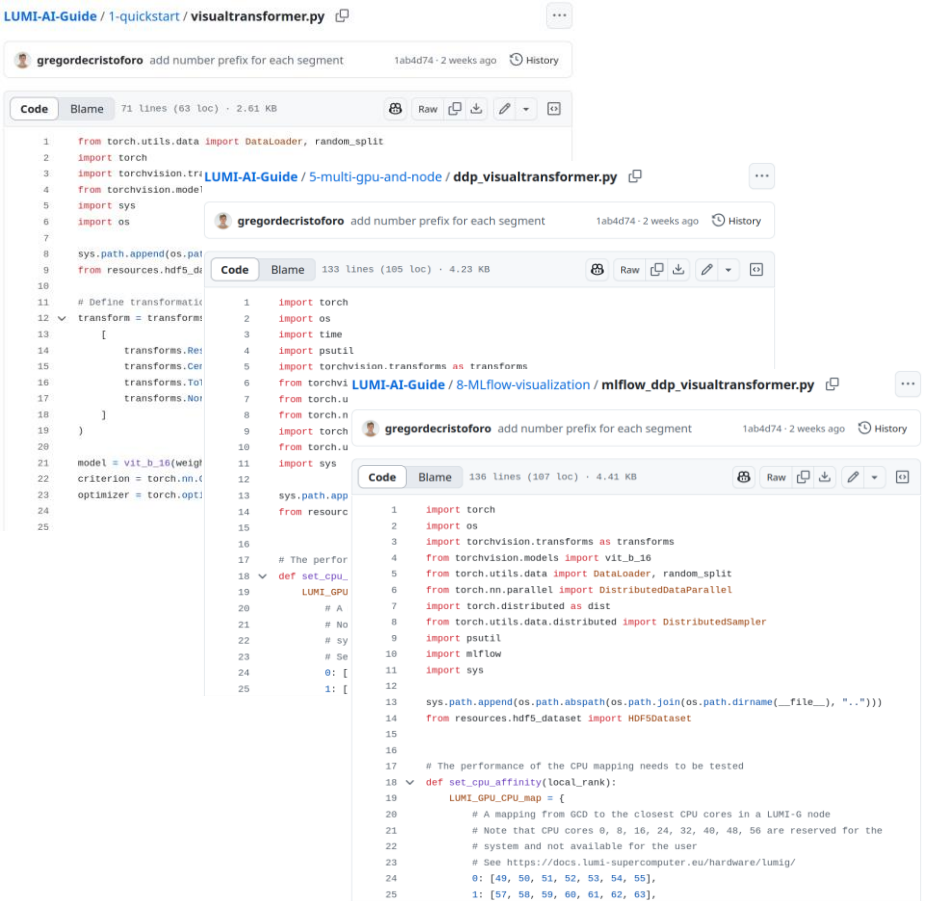

Please provide feedback

LUMI

- Find guide at https://github.com/Lumi-supercomputer/LUMI-AI-Guide
- Reach out to us via tickets or GitHub issues
- Is an essential part missing in the Guide?
- AI Guide will be regularly updated and maybe extended based on your feedback