2–3 May 2022
Hotel Europa, Bologna

The ML_INFN initiative

L. Anderlini [1], *on behalf of the ML_INFN initiative*

Istituto Nazionale di Fisica Nucleare
FIRENZE

Machine Learning Technologies for INFN

Most of the experiments and initiatives produce, analyse or process digital data.

Enthusiasm on the modern data processing technologies!

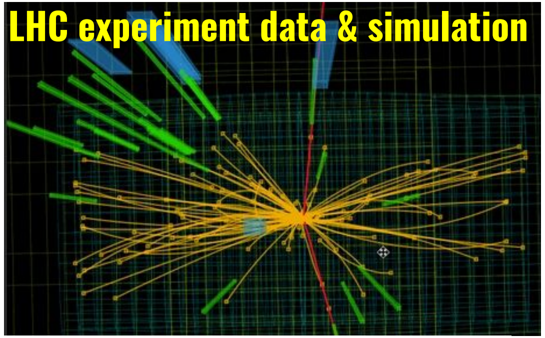

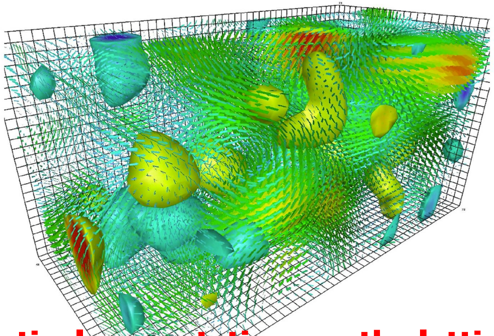

Gravitational wave detection

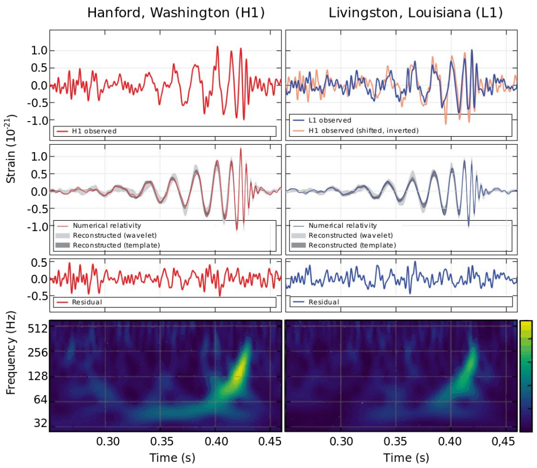

Raw radiation detector data

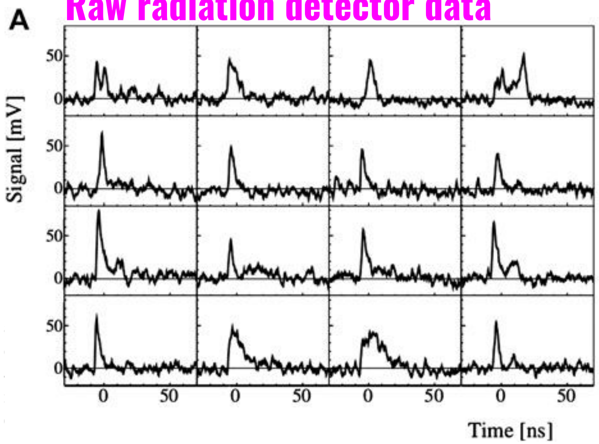

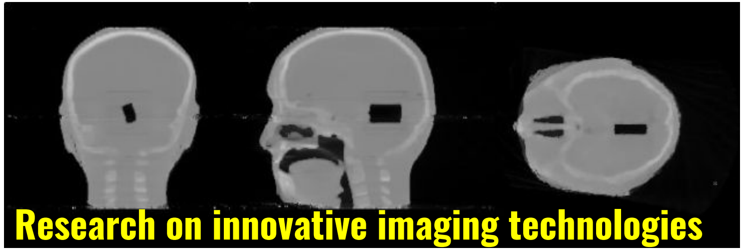

The potential barriers

Employing machine learning techniques often requires:

- specialized hardware and software setup
- specific training to identify tools and learning resources
- a community of experts providing support to research use cases

Lowering The potential barriers with ML_INFN

Employing machine learning techniques often requires:

- specialized hardware and software setup

WP1: provide a centrally maintained cloud-based infrastructure for interactive and batch ML fast prototyping, with access to modern GPU hardware and systems tuned for ML performance

- specific training to identify tools and learning resources

WP2: organize national training events for INFN users (Machine Learning hackathons)

- a community of experts providing support to research use cases

WP3: provide and organize example applications in a knowledge base

The numbers of ML_INFN

11 INFN **structures** involved in the preparation of the knowledge base

79 **researchers** devoting a fraction of their time to promote ML techniques for research

16 professional **GPUs** made available and accessible through the INFN Cloud Interface

110 **participants** to the **hackathons**, ranging from students to permanent staff members

WP1. The infrastructure

INFN Cloud

ML_INFN is built on top of **INFN Cloud**: a data lake-centric, heterogeneous federated Cloud infrastructure spanning multiple sites across Italy, providing an extensible portfolio of solutions tailored to **multidisciplinary scientific communities**.

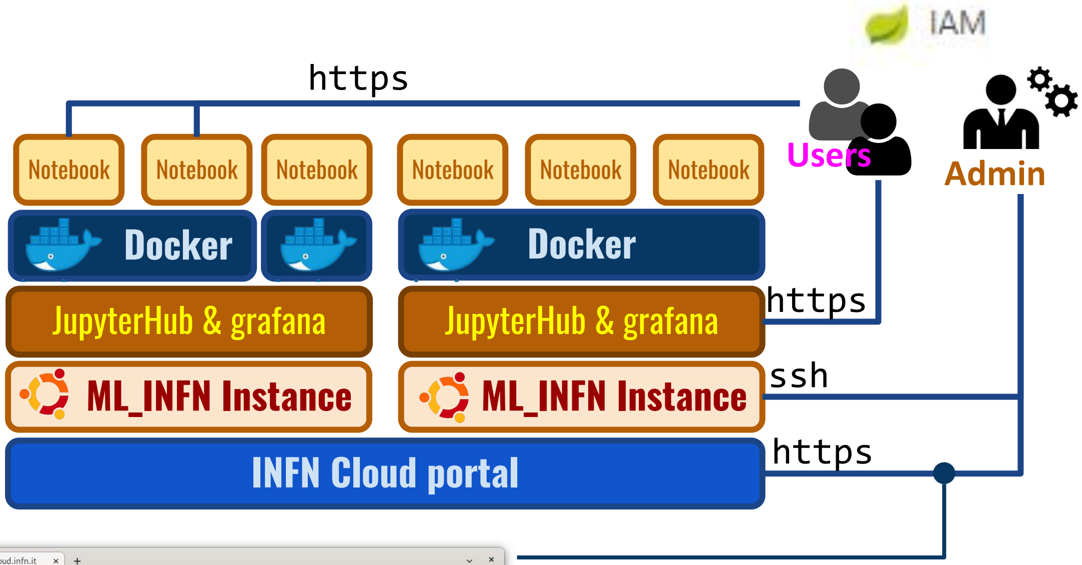

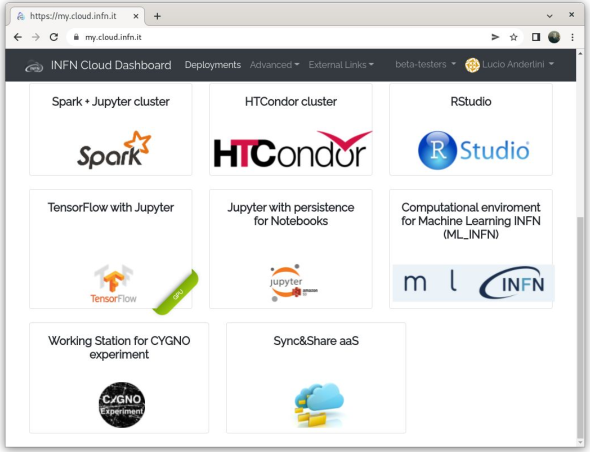

Federated baremetal resources

1× SuperMicro + 1× E4 servers:


- 1 TB RAM
- 64-128 CPU cores
- 36 TB local storage (NVMe)
- 8× **Tesla T4** GPUs
- 5× **RTX 5000** GPUs
- 1× **A30** GPU
- 2× **A100** GPUs
- 10 GbE connection to CNAF resources

Federated to CNAF OpenStack and INFN Cloud

Storage solutions

Storage from CERN experiments can be mounted with NFS from the Tier-1 storage

Hypervisors integrated to Ceph to manage persistent virtual volumes accessed from the VM with POSIX

WP2. Stewardship

Hackathons in Covid-19 Era

Originally planned as satellite events of scientific workshops, canceled due to pandemic, hackathons have been transformed in virtual events.

Registrations limited to 60 to guarantee decent tutor-per-student and RAM-per-student ratios.

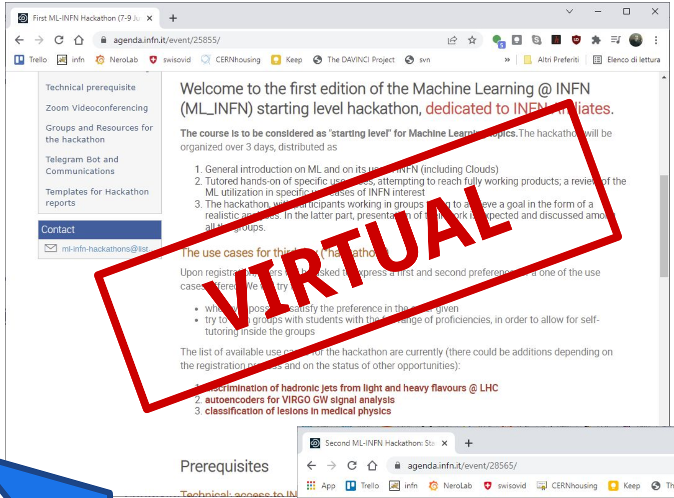

“sold-out” in 72 hours

Repeated

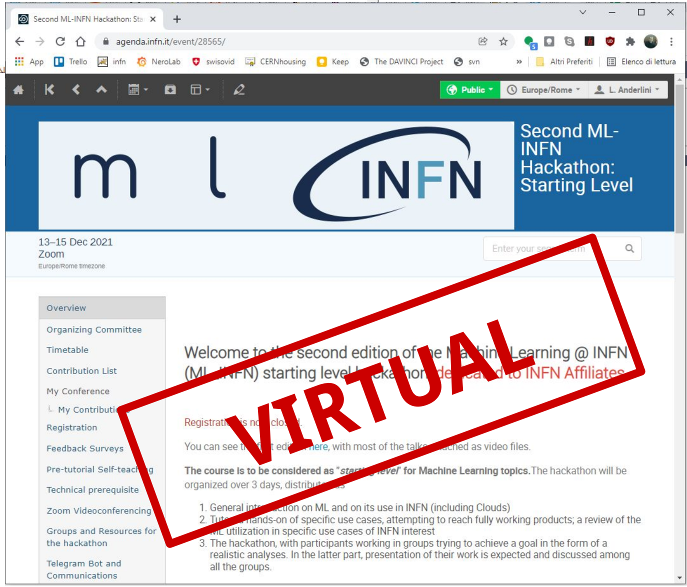

Lecture Program

Day 1	Day 2	Day 3
Lectures **Theoretical introduction to ML**	Hands-on **Neural Networks**	Hackathon **Exercises with tutors *continuous support***
Lectures **Cloud and Cloud Resources**	Seminars **Deep Neural Network Applications to INFN research**	
Lunch break		
Hands-on **Numpy, Pandas and Keras**	Hands-on **Exercises with tutors *on demand***	Closure **Reports from the students**

Hackathon use cases: *10 groups, one tutor per group*

Jet b-tagging at CMS

Recurrent Neural Networks with LSTM

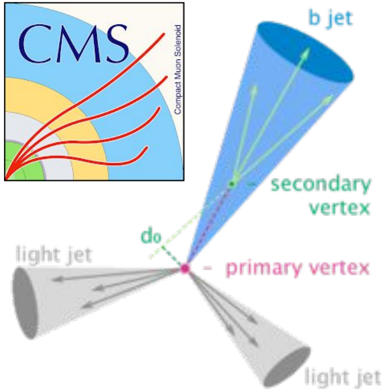

Higgs searches at CMS

Deep Neural Networks and Advanced Keras

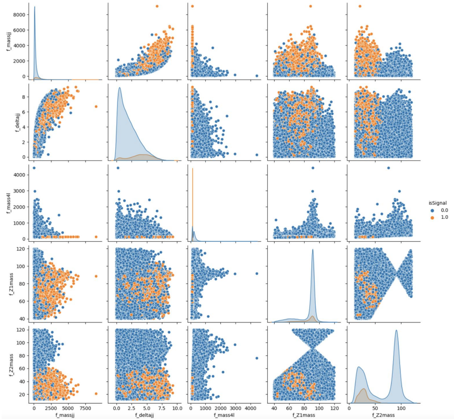

Gravitational Waves with Virgo

Autoencoders, anomaly detection and compression

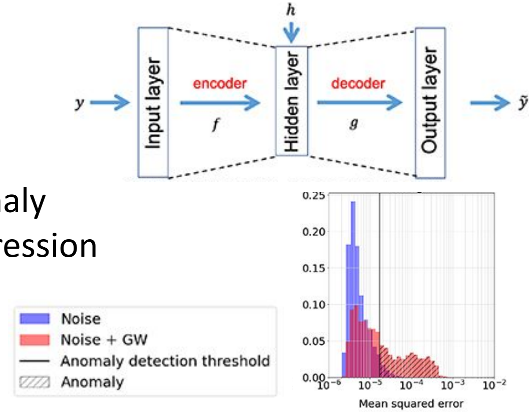

Segmentation of CT scans

Convolutional Neural Networks Handling 2D and 3D datasets

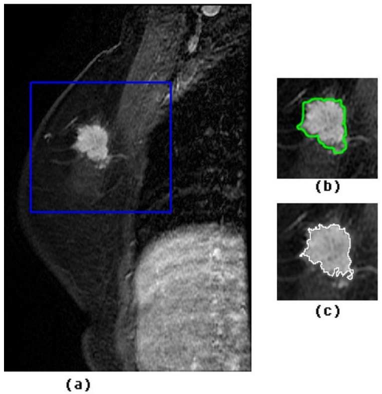

Final survey

A satisfaction questionnaire was submitted to the participant at the end of each event.

Generous feedback on:

- Level of difficulty
- Relevance and interest
- Technical setup

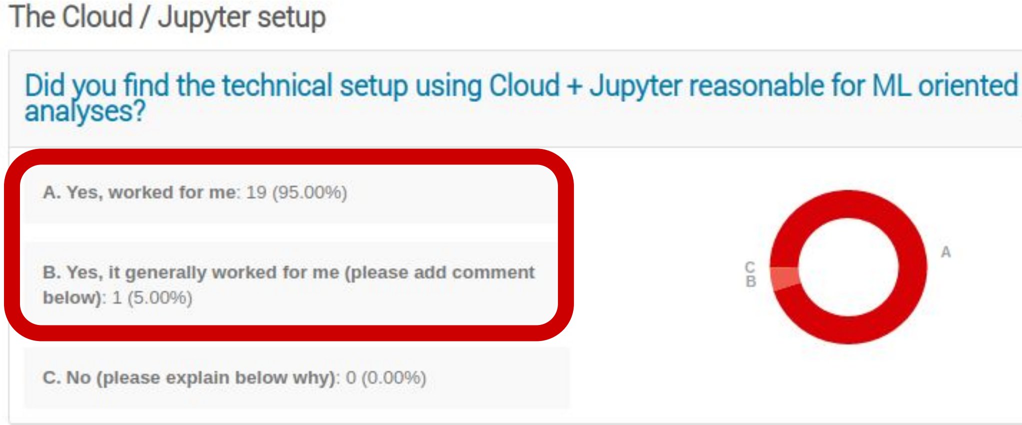

On the difficulty level:

On the choice of topic:

About a half of the participant responded they would have been willing to spend **more time** on the hands on and/or the hackathon.

Advanced Hackathon in Bari

21 – 24 November 2022

Registrations for the upcoming Advanced Hackathon are open. Register now.

Planned exercise topics:

- Image segmentation in Medical Physics
- Domain Adaptation in HEP
- Graph Neural Networks and Transformers
- Explainable AI

The event will be *in person* and a maximum of 20 participants will be accepted.

Objectives of the initiative:

- ➢ Present and discuss **realistic applications** of advanced algorithms to INFN research, **looking into the code**;
- ➢ Advertise **INFN Cloud** for sharing computing resources;
- ➢ Enhance **networking** of advanced ML practitioners within INFN.

WP3. Network and Knowledge Base

Confluence Knowledge Base

Atlassian Confluence was used to build a **Knowledge Base** reporting several machine-learning use cases, including those discussed at the hackathon.

Each entry includes:

- Runnable **example** as a jupyter notebook or a git repository
- **Contact information** of one or more experts

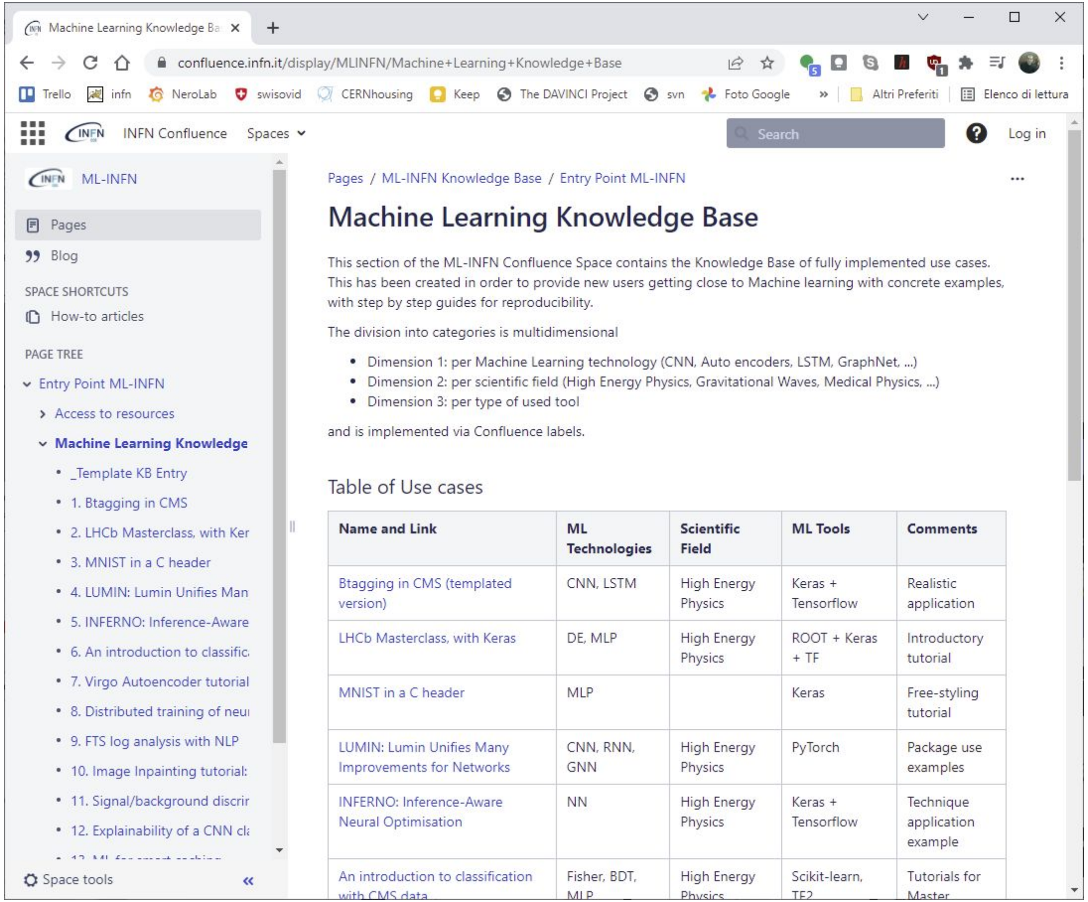

Publications

A. Abba et al., “The novel Mechanical Ventilator Milano for the COVID-19 pandemic featured”, Physics of Fluids 33, 037122 (2021)

L. Banchi *et al., “*Measuring Analytic Gradients of General Quantum Evolution with the Stochastic Parameter Shift Rule”, Quantum 5, 386 (2021).

L. Banchi *et al.* “Generalization in Quantum Machine Learning: A Quantum Information Standpoint” , PRX Quantum 2, 040321

P. Braccia *et al.,* “How to enhance quantum generative adversarial learning of noisy information”, New J. Phys. 23 053024 (2021)

D. Carlotti *et al.,* “Deep learning method for TomoTherapy Hi-Art: prediction three-dimensional dose distribution”, RADIOTHER ONCOL 161 (2021)

S. Francescato *et al.* “Model compression and simplification pipelines for fast deep neural network inference in FPGAs in HEP”, Eur. Phys. J. C 81, 969 (2021).

G. Graziani et al., “A Neural-Network-defined Gaussian Mixture Model for particle identification applied to the LHCb fixed-target programme”, JINST 17 (2022) P02018

A. Palermo *et al. “Machine learning approaches to the QCD transition”, LATTICE 2021, arXiv:2111.05216*

… and counting

BEST PAPER AWARD

ICCSA 2022 – International Conference on Computational Science and Its Applications

presented to paper titled

Hyperparameter optimisation of Artificial Intelligence for digital REStoration of Cultural Heritages (AIRES-CH) models

authored by:

Alessandro Bombini - Istituto Nazionale di Fisica Nucleare, Florence Sec.
Lucio Anderlini - Istituto Nazionale di Fisica Nucleare, Florence Sec.
Luca dell'Agnello - Istituto Nazionale di Fisica Nucleare, CNAF sec.
Francesco Giacomini - Istituto Nazionale di Fisica Nucleare, CNAF sec.
Chiara Ruberto - Florence University & Istituto Nazionale di Fisica Nucleare
Francesco Taccetti - Istituto Nazionale di Fisica Nucleare, Florence Sec.

In recognition of your outstanding contribution in the 22th International Conference on Computational Science and Applications (ICCSA 2022) and its associated workshops, held in collaboration with the University of Malaga, Spain, July 4 - 7, 2022.

ICCSA 2022 General Chairs
Osvaldo Gervasi, University of Perugia, Italy
Eligius Hendrix, University of Malaga, Spain
Bernady O. Apduhan, Kyushu Sangyo University, Japan

ICCSA 2022 Program Chairs
Beniamino Murgante, University of Basilicata, Italy
Inmaculada Garcia Fernandez, University of Malaga, Spain
Ana Maria A.C. Rocha, University of Minho, Portugal
David Taniar, Monash University, Australia

ICCSA 2022 Workshop and Session Organizing Chairs
Beniamino Murgante, University of Basilicata, Italy
Chiara Garau, University of Cagliari, Italy
Sanjay Misra, Ostfold University, Halden, Norway

ICCSA 2022 Award Chair
Wenny Rahayu, La Trobe University, Australia

Summary and conclusions

Summary

The ML_INFN initiative has been providing many INFN experiments with the hardware and the knowledge base to assess the potential **benefit of machine learning to their research** for three years.

The **ML_INFN** project relies on **INFN Cloud** solutions and it federates resources optimized for ML performance in interactive and batch-like usage patterns (high-end professional GPUs, NVMe disks, many-core high-RAM systems)

A series of national training events (***machine learning hackathons***) and a collection of tutorials and real applications within the INFN community (***knowledge base***) contribute to building **a network of experienced and enthusiast machine learning practitioners**, lowering the skill gap to benefit from machine learning developments.

Outlook

Machine Learning is here to stay. In the next future:

- We will organize ***Advanced Training Course(s)*** on Deep Learning
- We will provide Cloud-based access to **FPGAs** as **Machine Learning accelerators**
  - Two U50 and a U250 Xilinx FPGAs recently federated to the cloud
  - Already accessed through virtual machines, aiming at provisioning FPGAs as a Service
- We will keep supporting students and researchers employing **Machine Learning technologies** in their daily activities.