CINECA

MLWP

AI applied to meteorological research

Matteo Angelinelli

30 May 2024

NOT-FOR-PROFIT CONSORTIUM

SINCE 1969 CINECA SUPPORTS THE ITALIAN ACADEMIC SYSTEM

112 MEMBERS

2 Ministries, 69 Universities, 41 Academic and Research Institutions

5 OFFICES

Bologna, Milan, Rome, Naples, Chieti

≃1000

Employees

≃ 110 MLN €

Yearly Revenue in 2021

50+ YEARS OF IT SERVICES

1969

SUPERCOMPUTING

4 founding Universities
Bologna, Florence, Padua, Venice

'80

MINISTRIES

IT Systems for the Italian Ministry of Universities and Research

'90

UNIVERSITIES

IT Systems for the Italian Academic System

2000

PA & INDUSTRY

Technologial Transfer to Healthcare
Public Administration
Industry

2020

HPC PRE-EXASCALE

Artificial Intelligence
Big Data
Quantum Computing

∞

HPC DATA ANALYTICS TEAM

OUR GOALS

Support researchers, data scientists and organisations in using high-performance computing resources to advance data analysis and artificial intelligence projects.

OUR SERVICES

Access to computing infrastructure
Access to data
PoC development
Consulting
Training

OUR DOMAINS

Industry 4.0
Meteorology and Climate
Cultural Heritage
Agrifood
Health and Society

AI APPLIED TO METEOROLOGICAL RESEARCH

- **FROM NWP TO MLWP:** AI applied to weather forecasting
- **THE BASICS OF AI:** data and computing power
- **THE DOWNSCALING PROBLEM:** from global to local forecasts
- **WEATHER FORECASTING:** multi-framework approaches for Italian MLWP models
- **WHAT'S NEXT?:** from forecast to social impact

CINECA

FROM NWP TO MLWP

AI applied to weather forecasting

FROM NWP TO MLWP

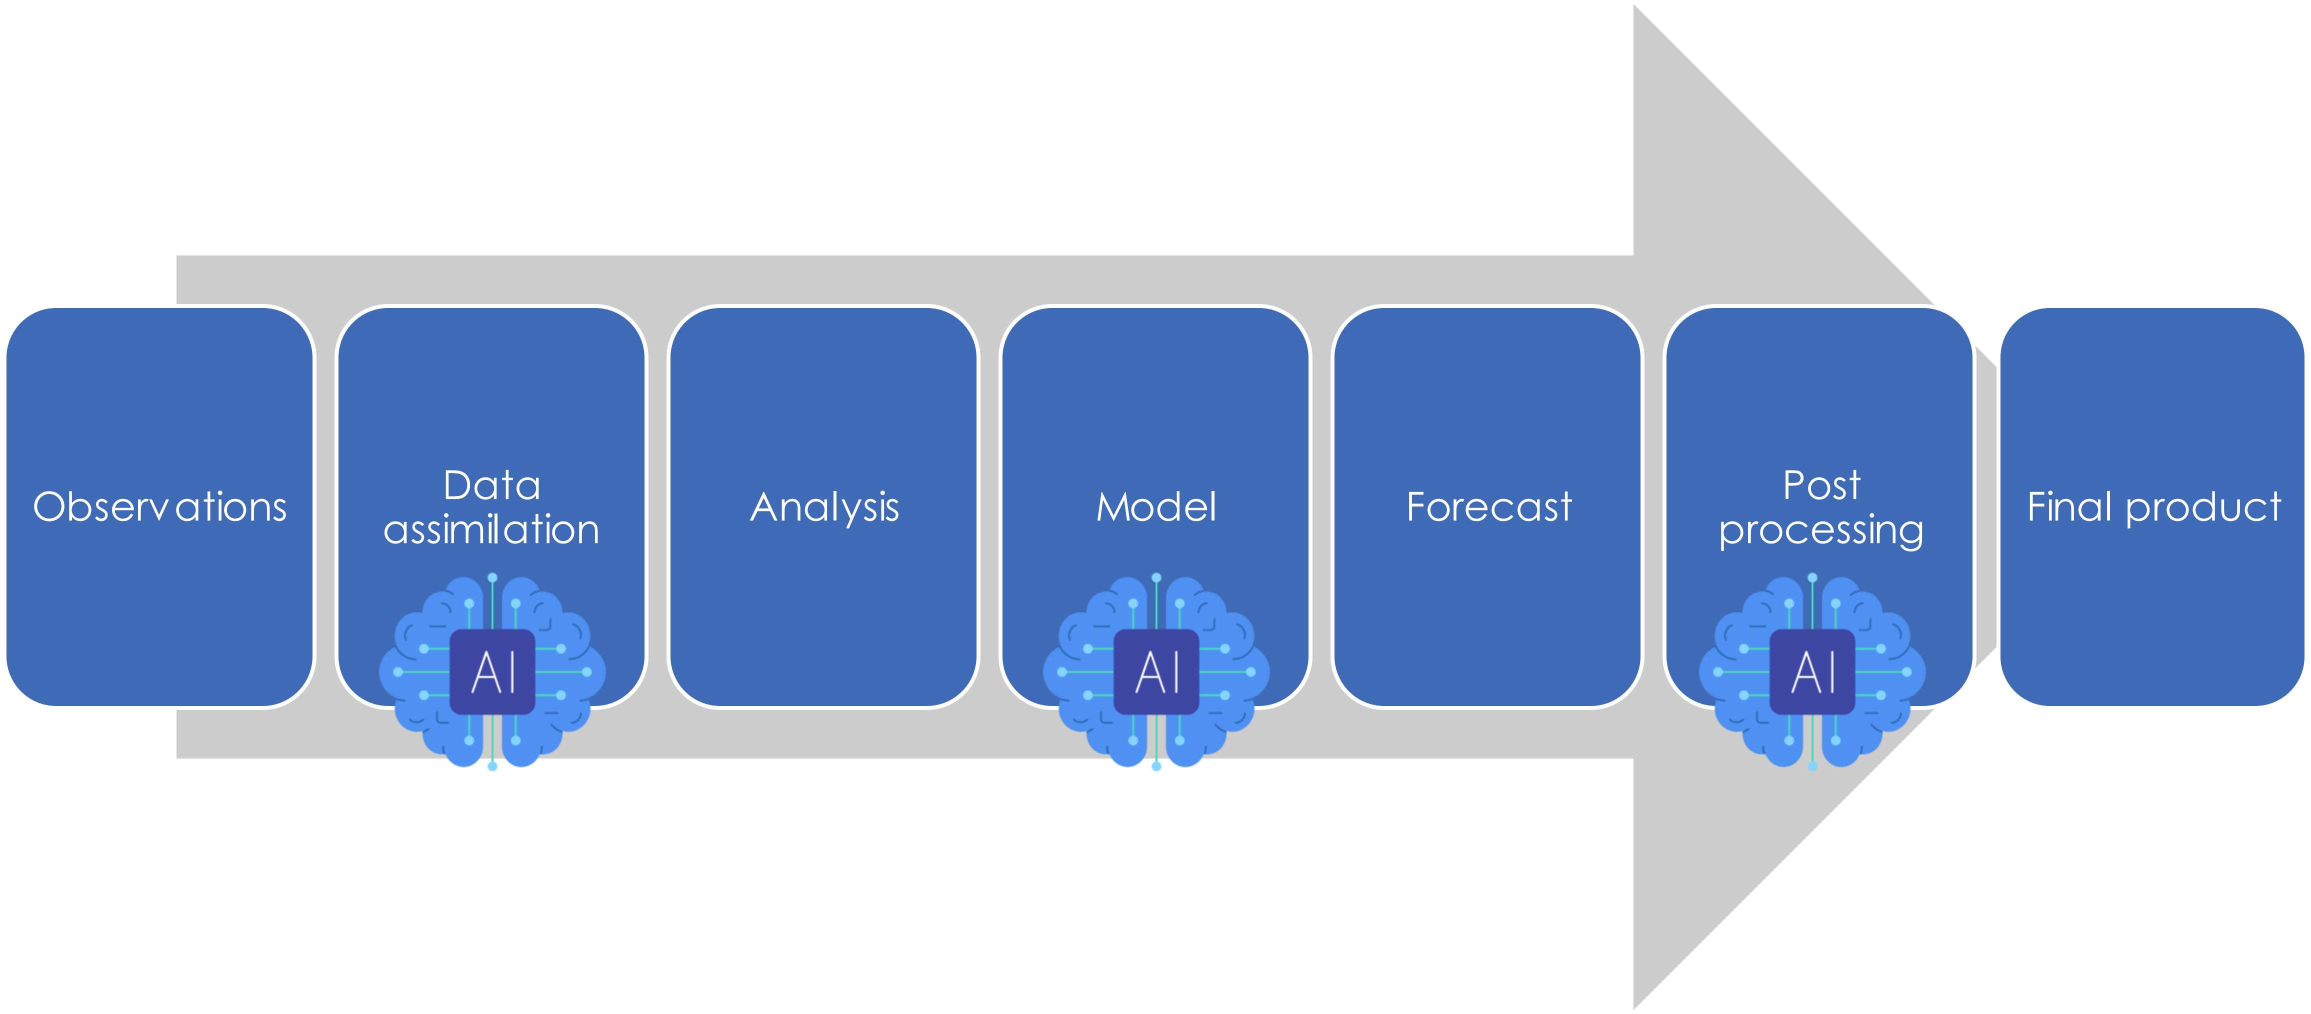

FROM NWP TO MLWP

Training Data	Approach	Inference input	Lead time	Resolution	Example
Observations	ML	Observations	Nowcasting (0-6 hrs)	Regional (sub-km/km)	GPTCast, LDCast
Observations Analysis	ML	Observations Analysis	Short Range (12-24 hrs)	Regional (km)	MetNet
Re-Analysis (ERA5)	ML	Analysis	Medium Range (10-15 days)	Global (0.25˚)	GraphCast, Pangu, FourCast
Re-Analysis (ERA5)	ML Physics	Analysis	Medium Range (10-15 days) Long Range (months/years)	Global (0.7˚)	NeuralGCM

FROM NWP TO MLWP

PROS	CONS
Accuracy	Strongly dependent on the re-analysis quality
Forecast computational resources	Resolution determined by the training dataset
Forecast computational time	Rain forecast
Open-source	Extreme events

PANGU-WEATHER

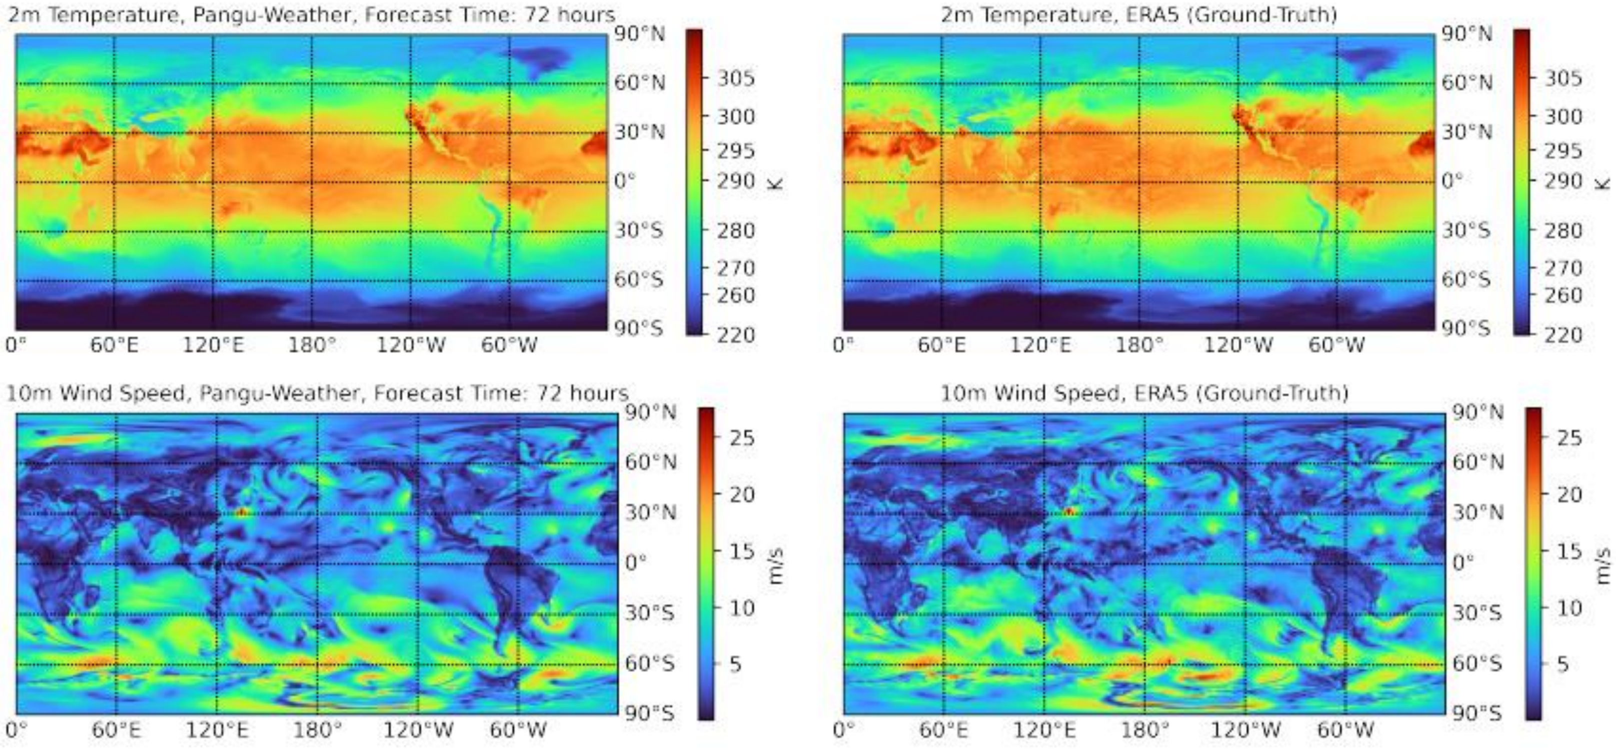

Developed by HUAWEI, PanguWeather uses a **Swin transformer** and a new algorithm based on the 3D Earth-specific transformer (3DEST), an encoder-decoder architecture derived from a variant of a vision transformer. This model features an Earth-specific positional bias that replicates the physical structure of the Earth. The input consists of 4 surface variables and 5 at 13 altitude levels, while the output includes 4 models with different lead times (1, 3, 6, and 24 hours). Training this model requires **approximately 16 days** on a cluster of 192 Nvidia Tesla-V100 GPUs.

GRAPHCAST

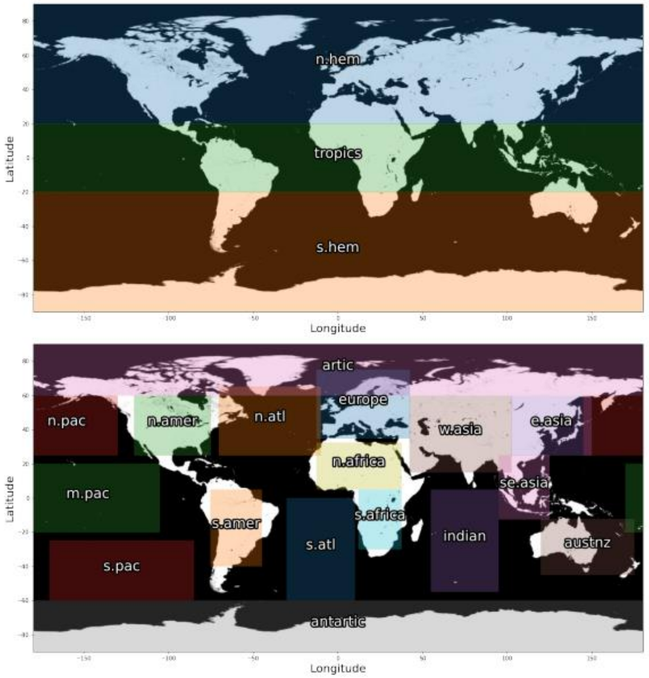

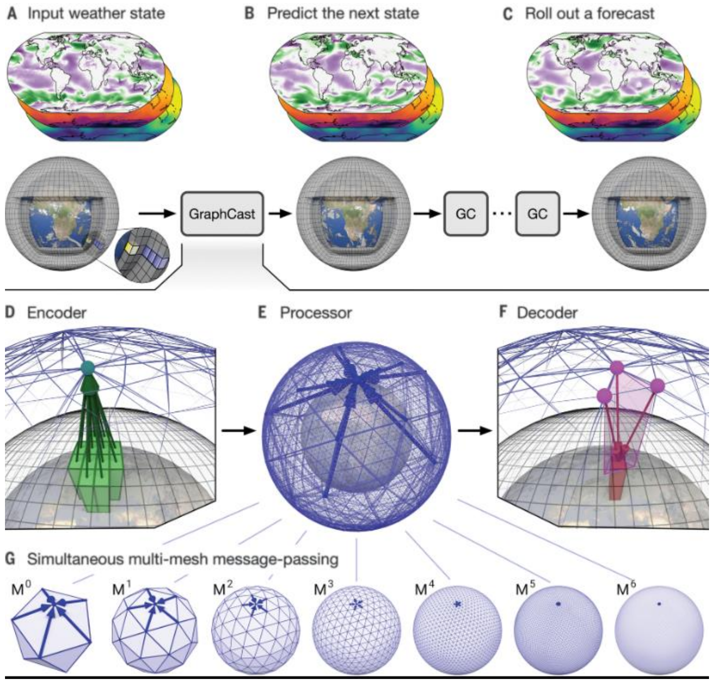

Developed by DeepMind (Google), this model is based on **graph representation**: the Graph Neural Network (GNN). It employs an encoder-process-decoder architecture with 36.7 million parameters. GNNs, coupled with fluids, have demonstrated remarkable capabilities in learning the dynamics of systems described by partial differential equations. The input comprises 5 surface variables, 6 at 37 altitude levels, observed every 6 hours. The output consists of a family of models utilizing two subsequent states of the atmosphere to generate forecasts. Training this model takes **approximately 4 weeks** on a cluster of 32 TPUs.

THE BASICS OF AI

Data and computing power

MISTRAL
Meteo Italian Supercomputing Portal

The objective of the MISTRAL portal is to **facilitate and promote the re-use of datasets by the meteorological community and related communities to provide value-added services using HPC resources**, thus creating new business opportunities. The data will be available as gridded fields, probabilistic products (such as rainfall forecasts for flash flood prediction) or point time series derived from the Italian operational forecast modelling chain and post-processed fields (such as thunderstorm probability).

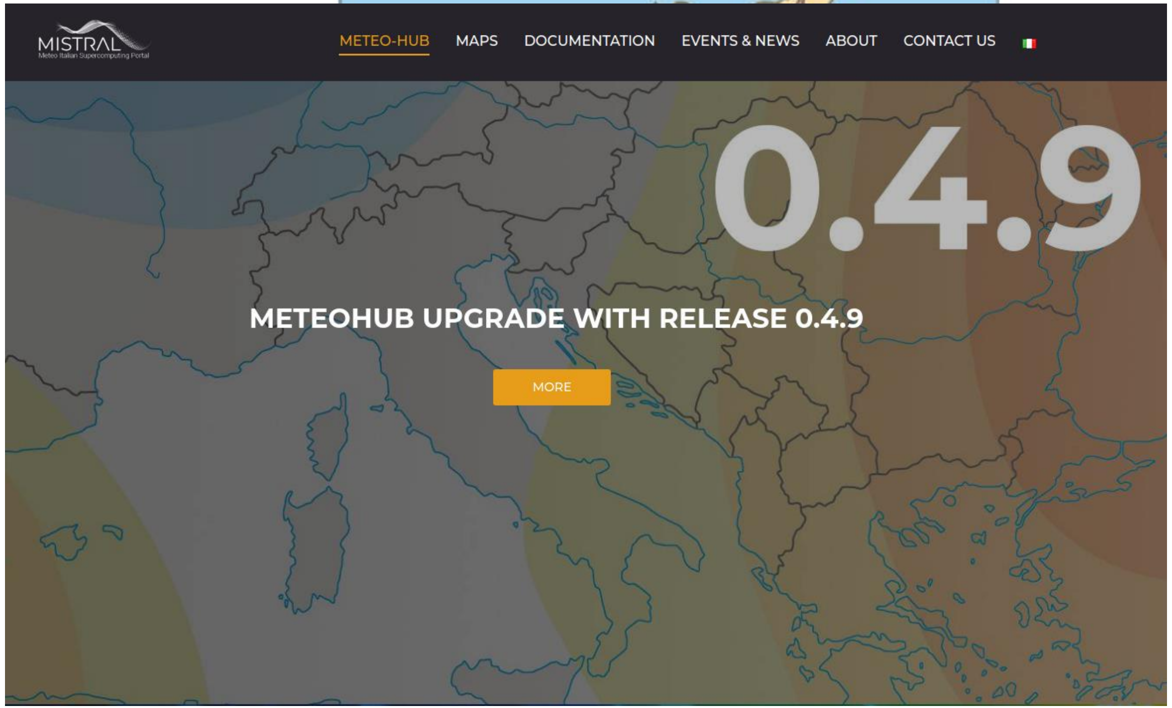

Highlander

High performance computing to support smart land services

Through the use of high performance computing, **Highlander will make it possible to process data for climate forecasting and reducing the risks associated with climate change,** for more intelligent and sustainable management of natural resources and the territory.

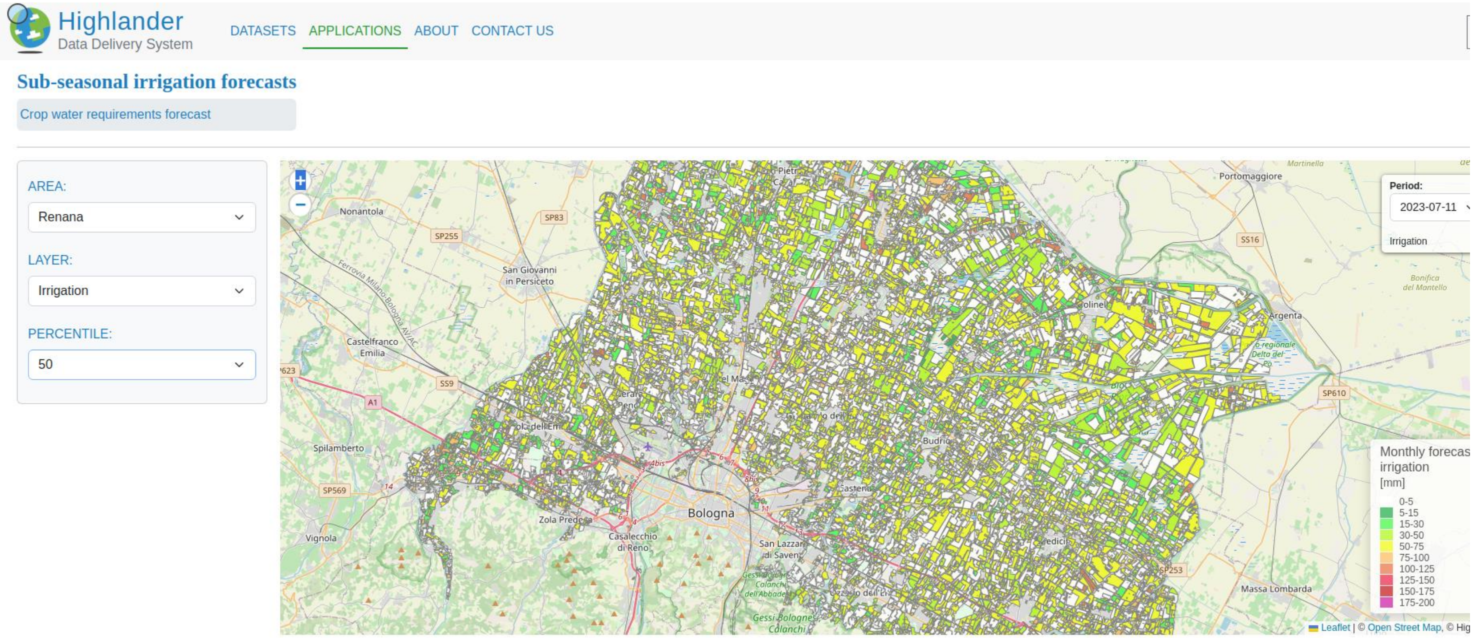

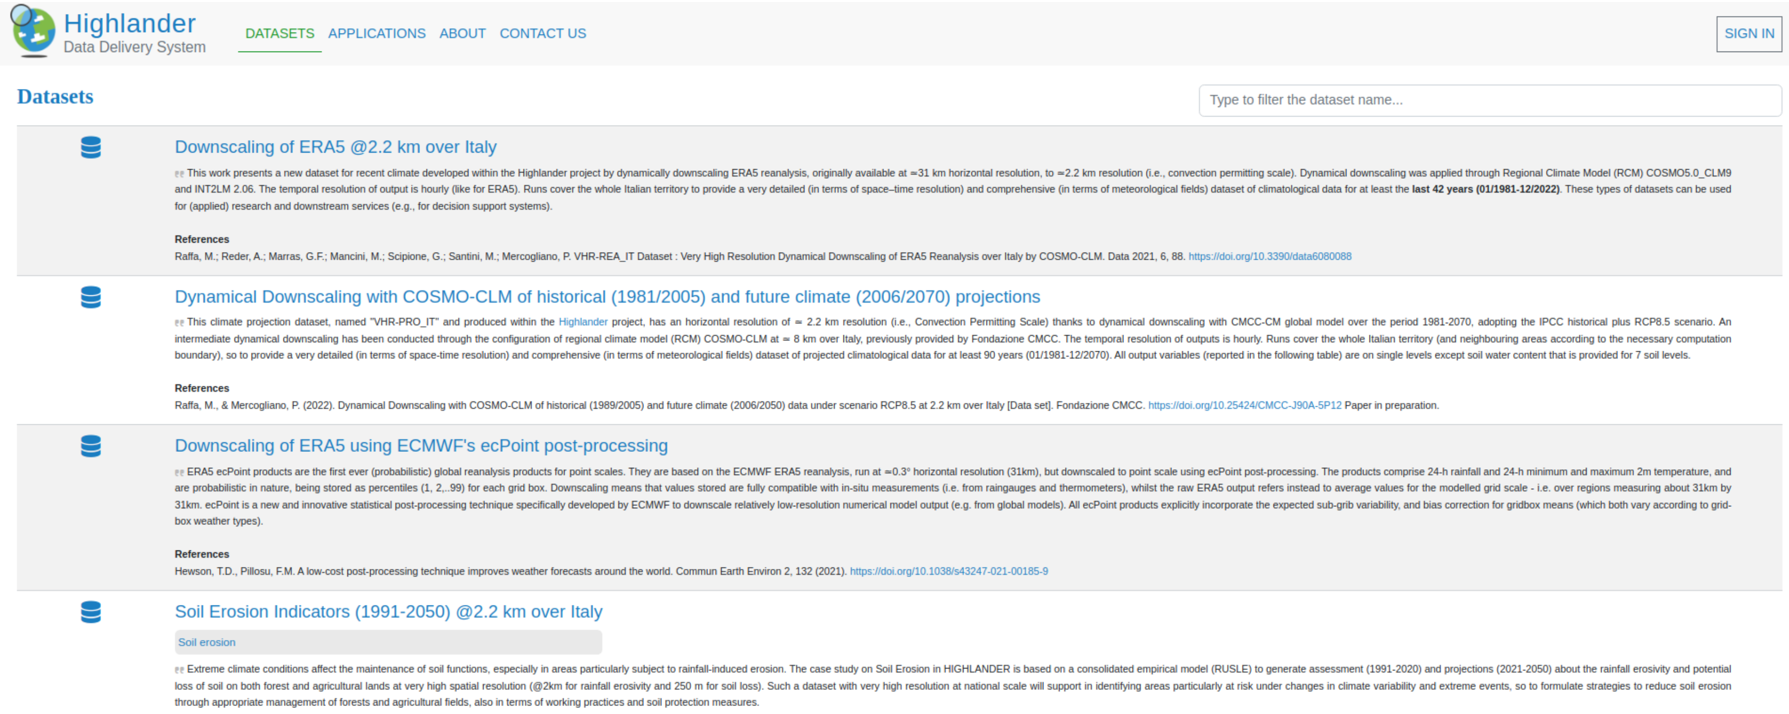

LEONARDO CINECA

Booster module

Nodes: 3456
Processors: Intel Xeon 8358 32 cores, 2.6 GHz
Cores: 110592 (32 cores/node)
Accelerators: 4XNvidia custom Ampere GPU 64GB HBM2
RAM: (8×64) GB DDR4 3200 MHz

General purpose module

Nodes: 1536
Processors: Intel Saphire Rapids 2×56 cores, 4.8 GHz
Cores: 172032 (112 cores/node)
RAM: (48×32) GB DDR5 4800 MHz

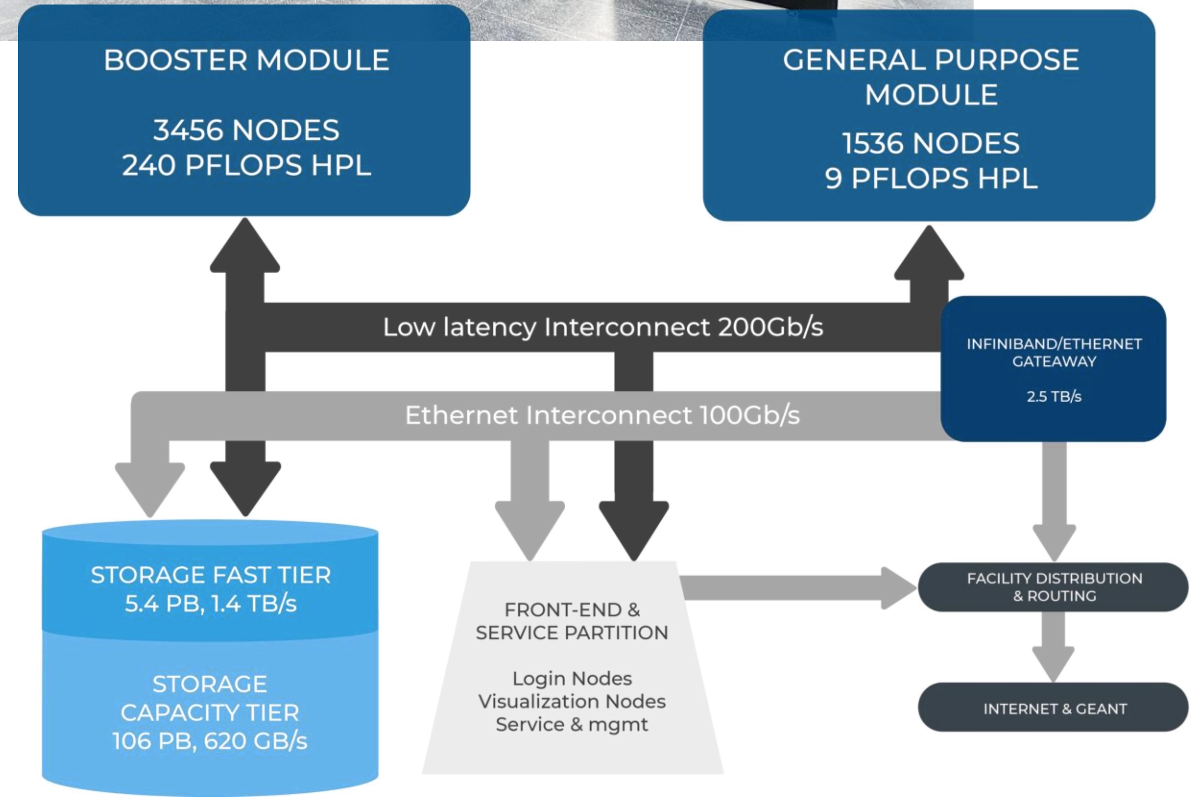

THE DOWNSCALING PROBLEM

From global to local forecasts

DOWNSCALING: FROM GLOBAL TO LOCAL FORECAST

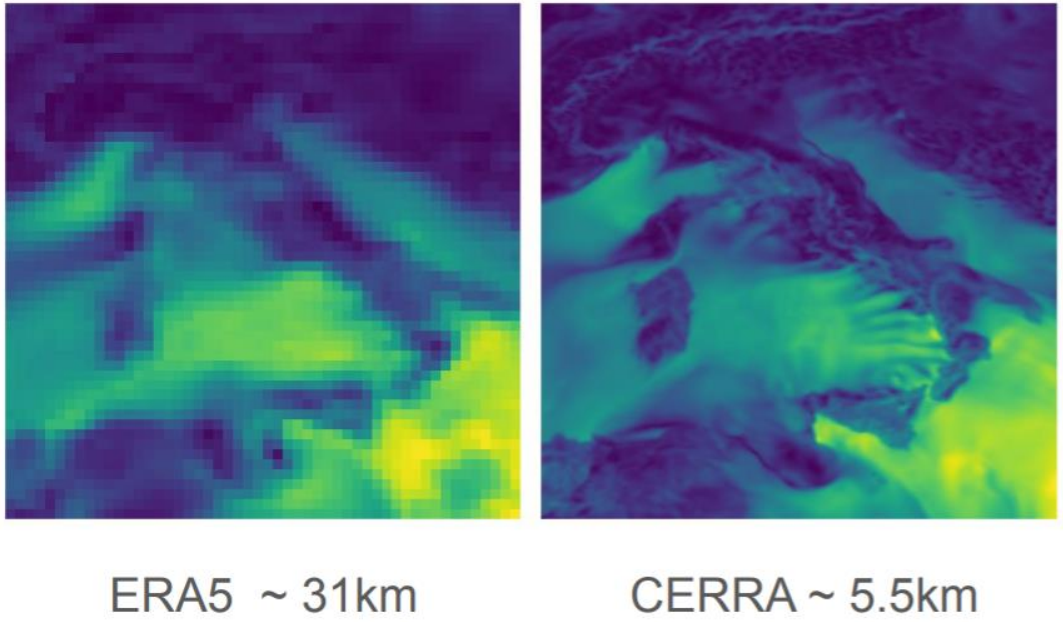

Downscaling in meteorology is a technique used to obtain **high-resolution weather forecasts** for specific areas by refining information from global or regional-scale weather models. Global weather models often have relatively low resolutions, typically spanning tens of kilometers, which may not capture local or detailed weather phenomena adequately. Downscaling involves employing high-resolution weather models to provide more detailed forecasts for smaller geographical areas. These models refine input data from global or regional models to produce **more precise predictions for small geographic areas**. Downscaling aims to enhance the accuracy and detail of weather forecasts for specific regions or areas of interest.

OUR DEEP LEARNING APPROACH FOR DOWNSCALING

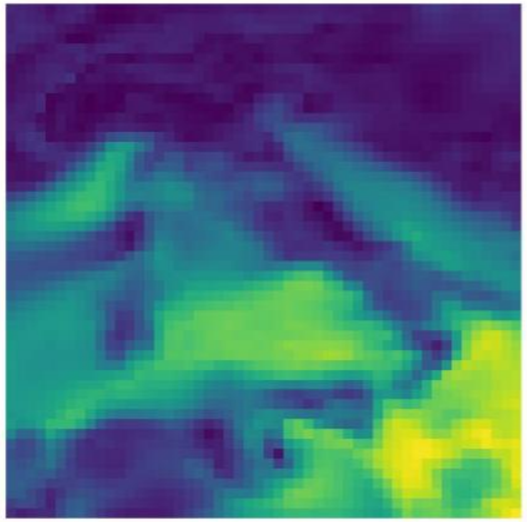

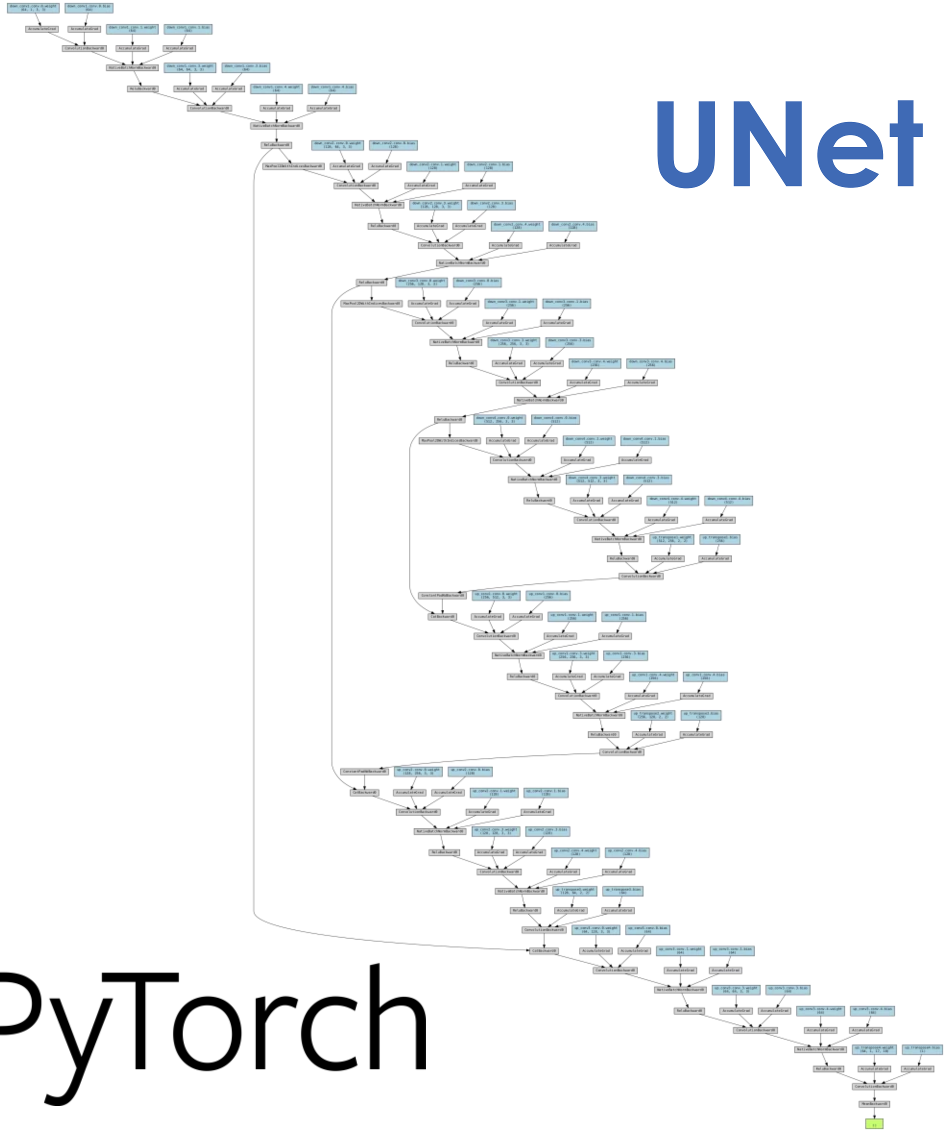

UNet

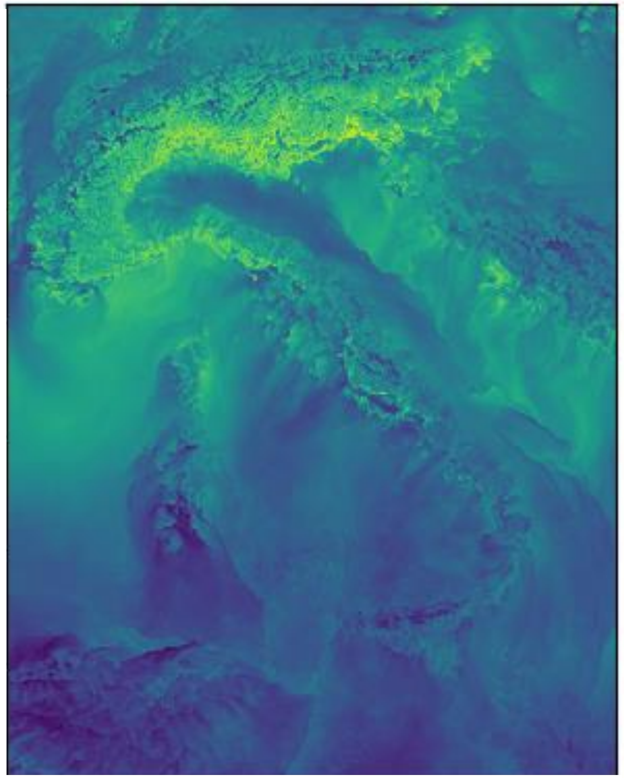

PyTorch

HOW WE USED PYTORCH FOR DOWNSCALING

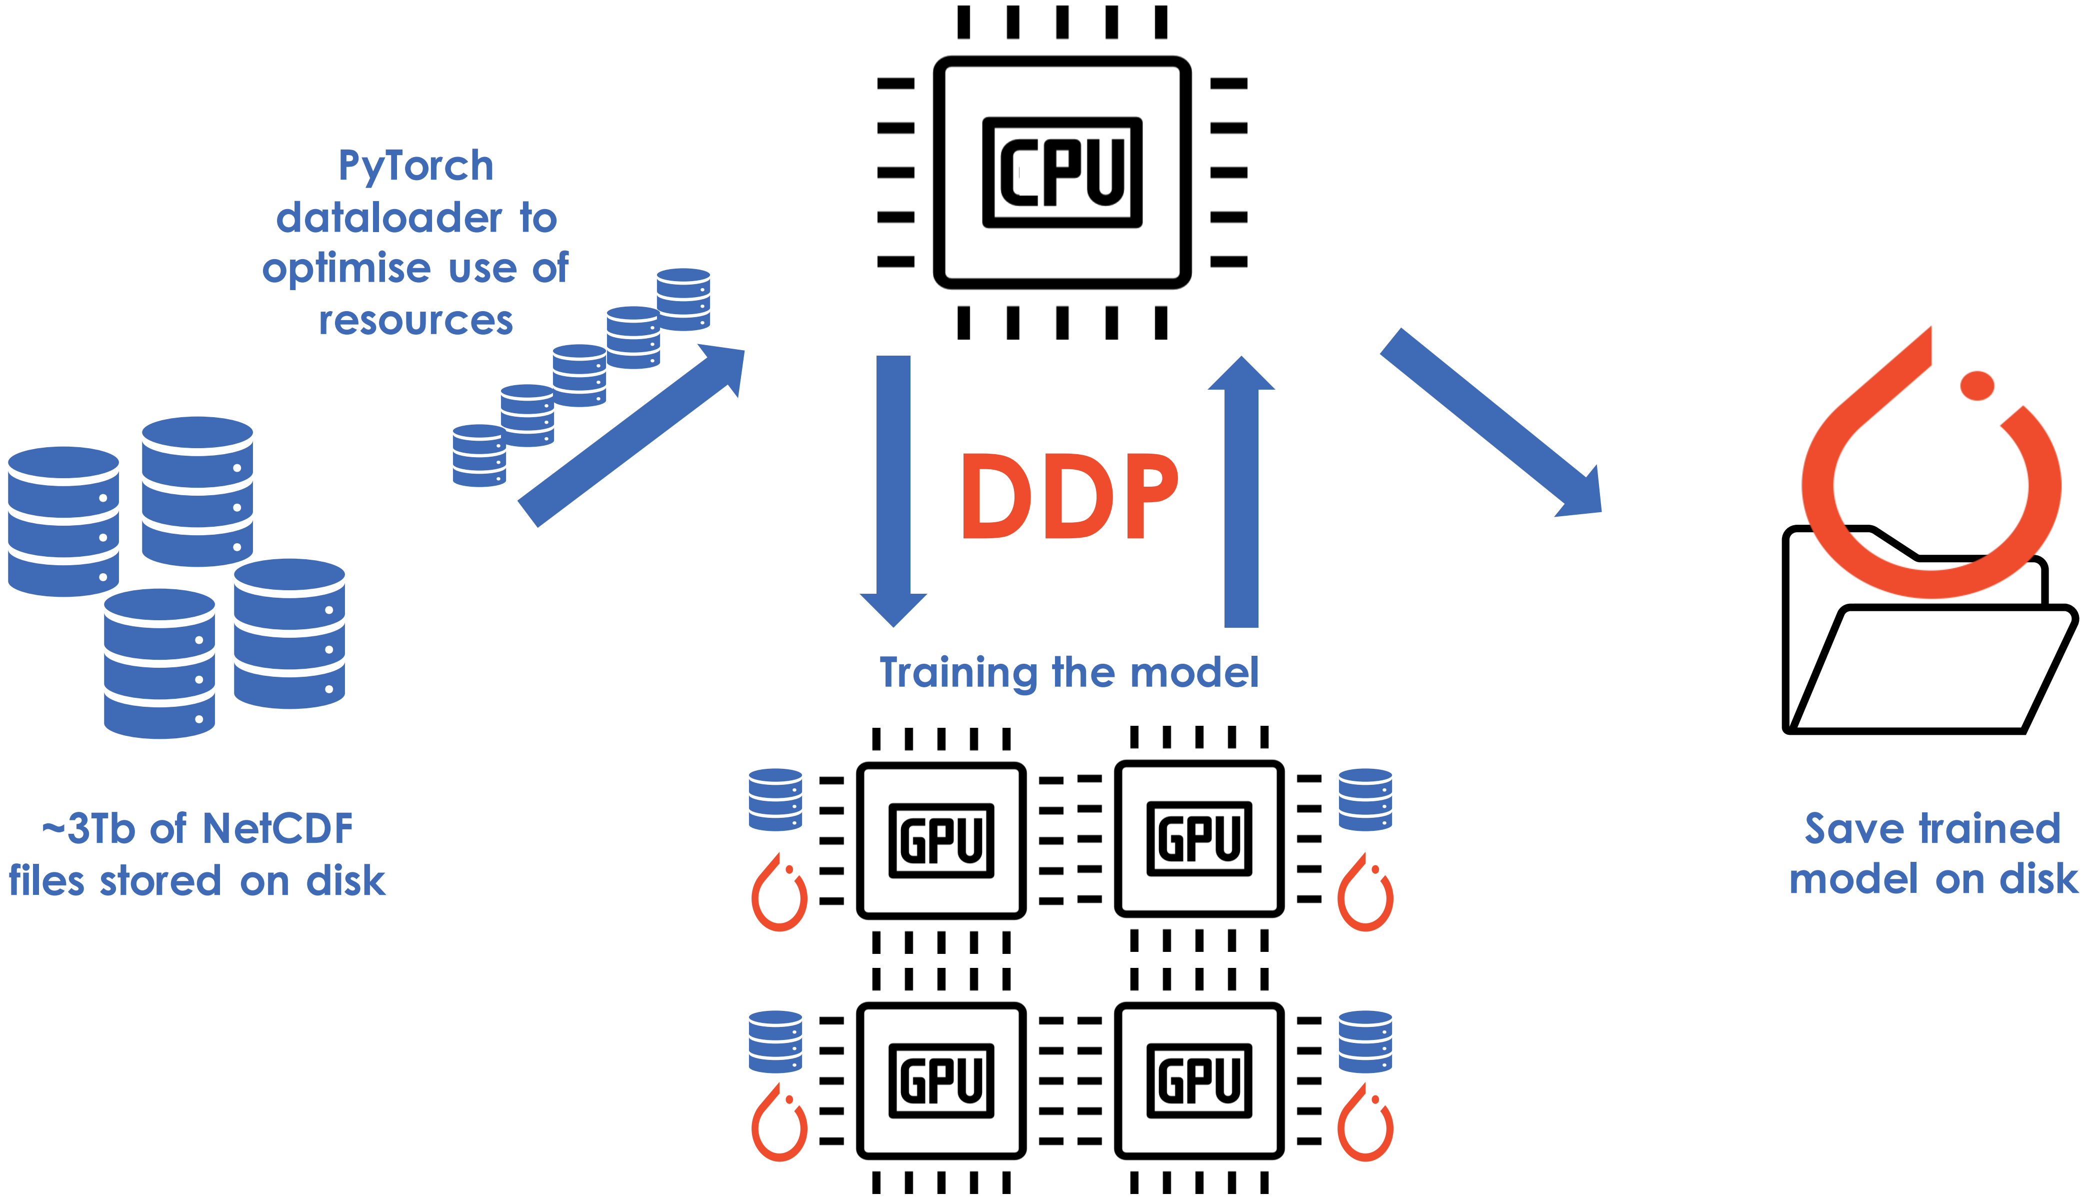

WEATHER FORECASTING

Multi-framework approaches for Italian MLWP models

AI-GCM

AI General Circulation Model

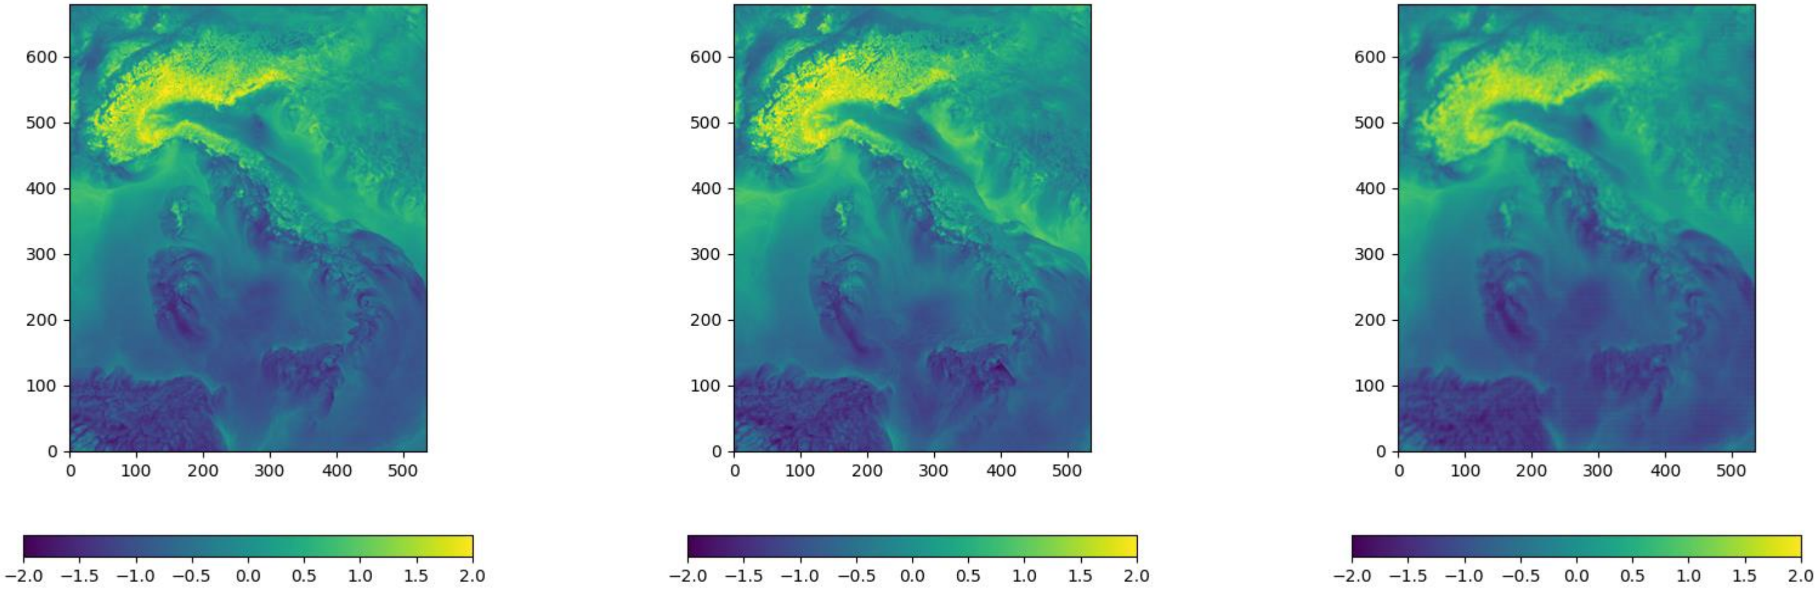

The project aims to construct a novel GCM approach **using recent developments in ML** and supported by intra-seasonal atmospheric signals. The aim is twofold. On the one hand, the potential to significantly reduce the run time of short-term forecasts. On the other hand, the possibility to increase the forecast skill for the medium term. The final objective is to produce a prototype model that could be further developed with future funding. The training process will use an ERA5 reanalysis dataset to develop different models based on **U-Net and Transformers technologies**. This project is funded by IFAB for 18 months. Cineca is involved by providing computational support and a suitable framework to run the model.

UNET FOR WEATHER FORECASTS

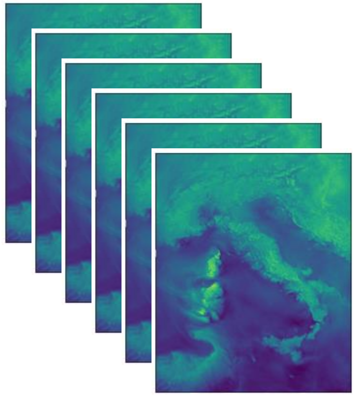

6 physical variable
at t=0

UNet

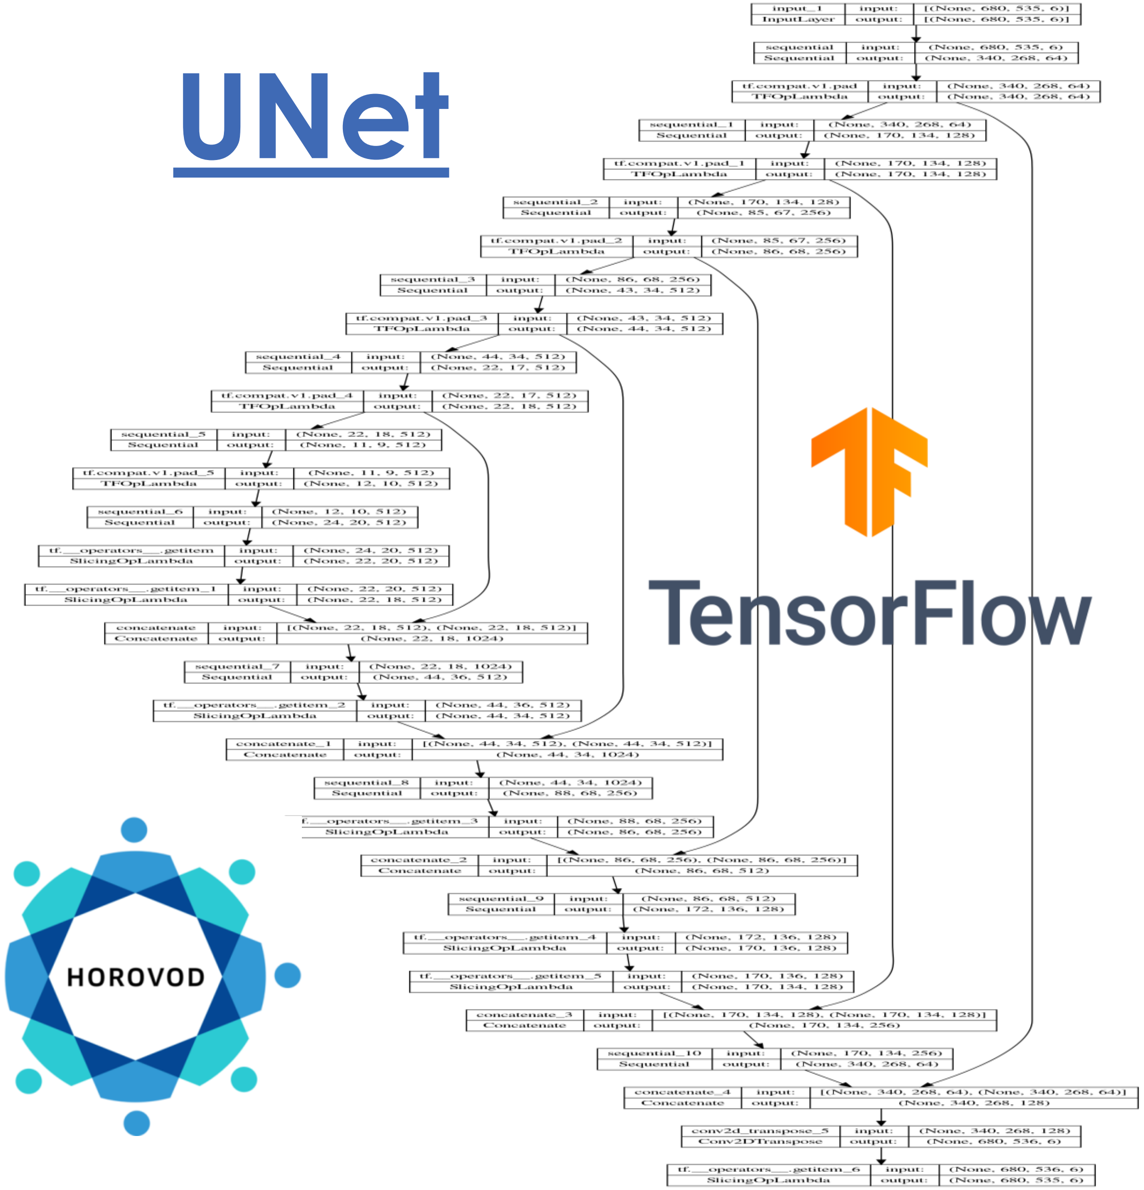

TensorFlow

HOROVOD

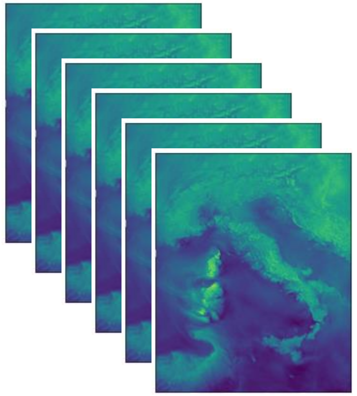

6 physical variable
at t=+3

OUR HPC IMPLEMENTATION

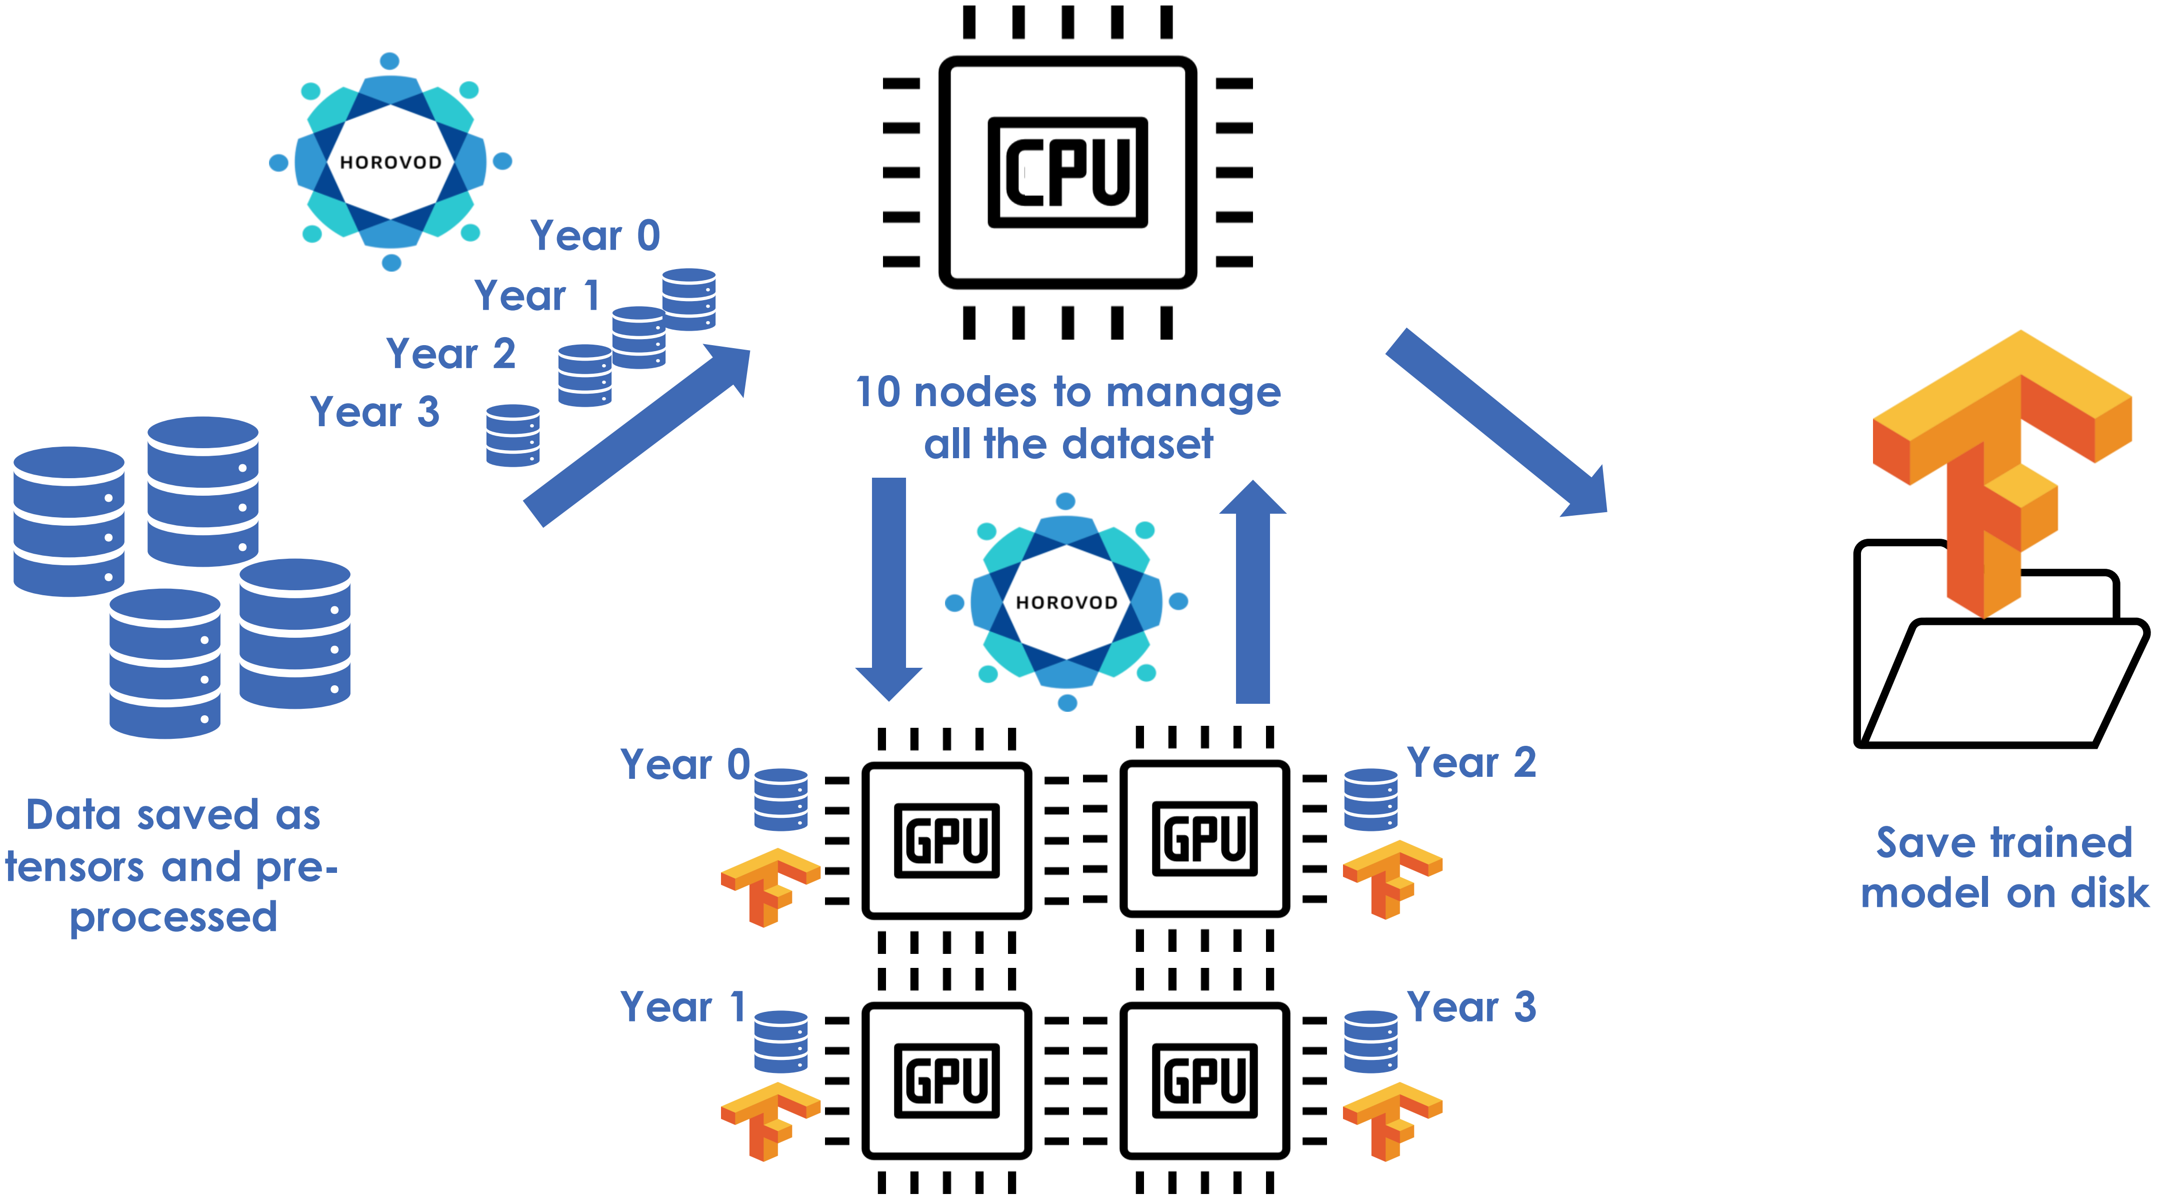

SPATIO-TEMPORAL: MOVIES MAKE PREDICTIONS

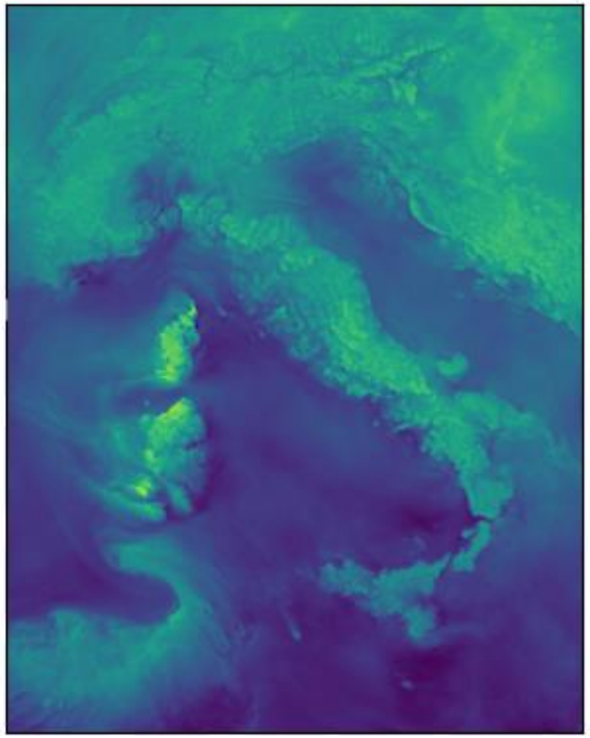

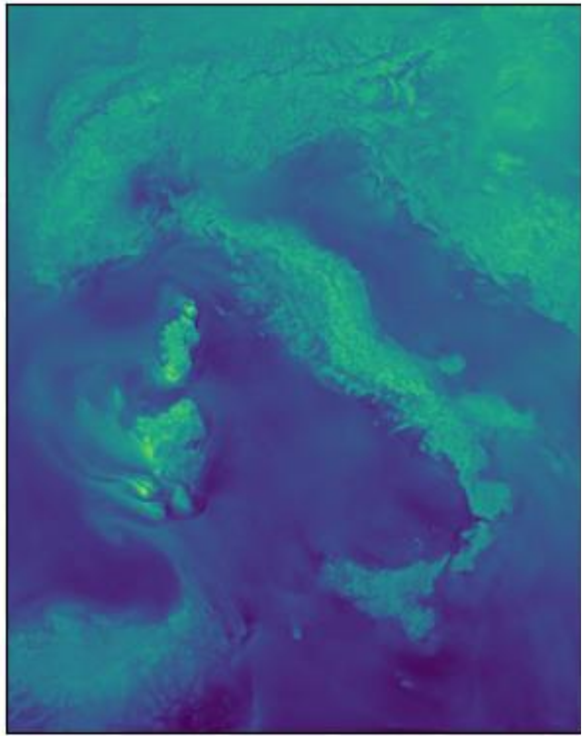

Spatio-temporal learning algorithms are designed to analyze and understand **spatial and temporal patterns in data**. They consider both the spatial and temporal dimensions of data, enabling them to capture spatial relationships between different entities and temporal dynamics over time. Typically leveraging convolutional neural networks or transformer, **they automatically learn these relationships from data**. Widely applied across diverse sectors including environmental sciences, geospatial data analysis, social network monitoring, and financial forecasting, they aim to comprehend and forecast spatio-temporal behaviors and trends within data.

ViT FOR WEATHER FORECASTS

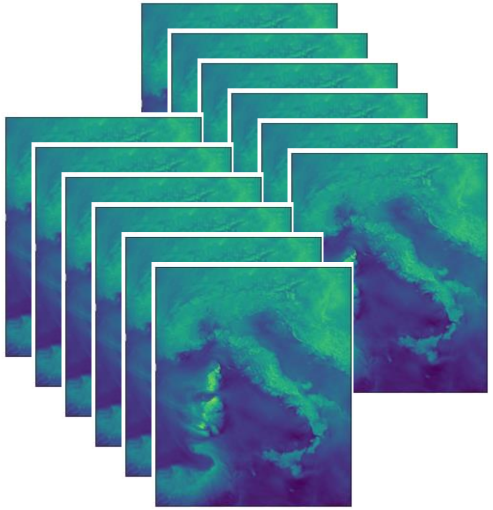

1 physical variable
from t=0 to t=12

OpenSTL: A Comprehensive Benchmark of Spatio-Temporal Predictive Learning

[OpenSTL paper]

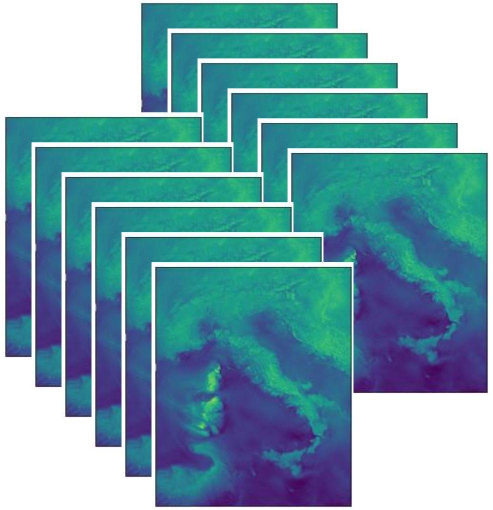

1 physical variable
from t=12 to t=24

OUR HPC IMPLEMENTATION

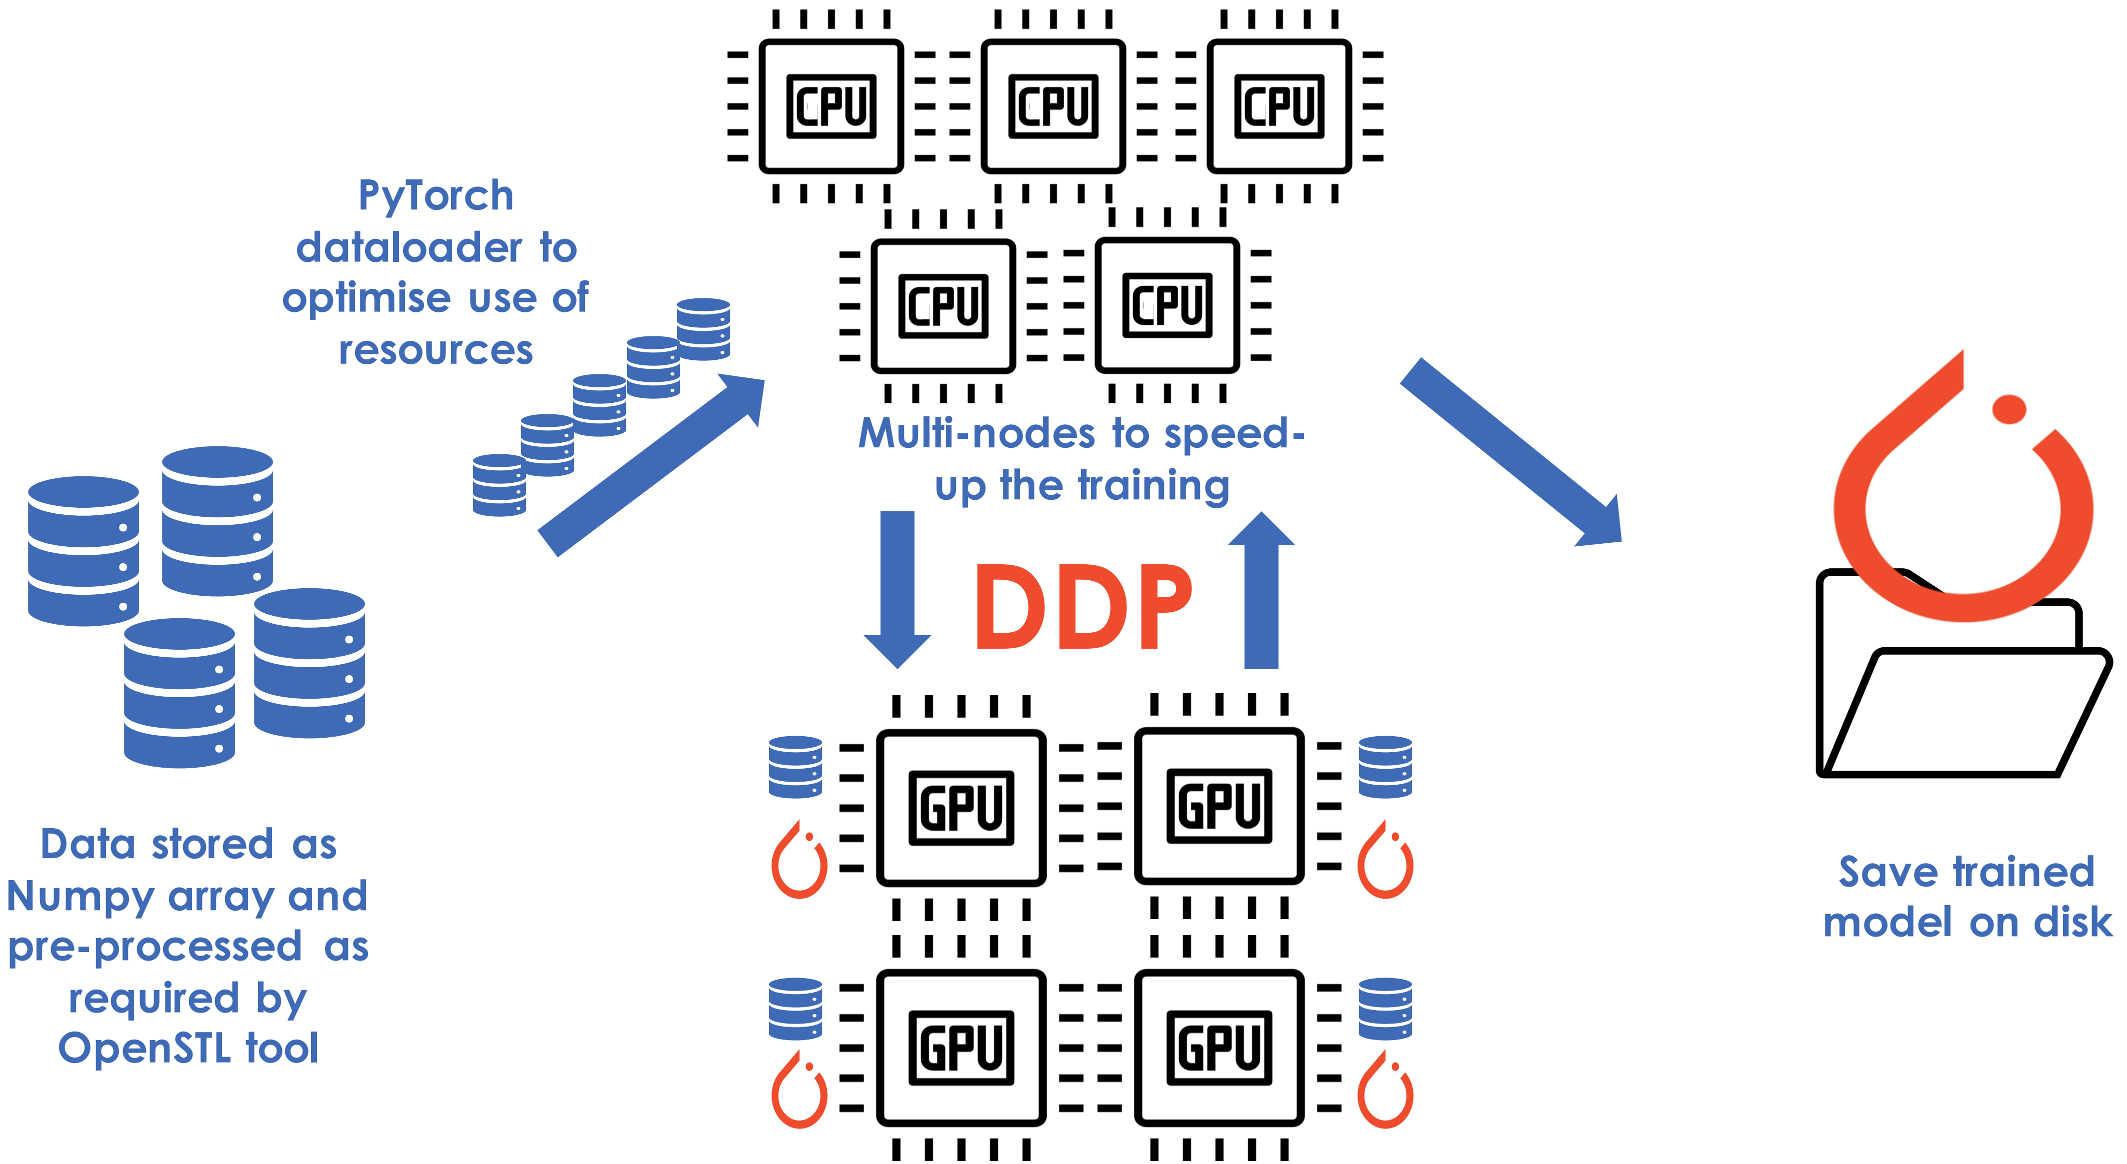

COMPARE, EVALUATE AND UNDERSTAND

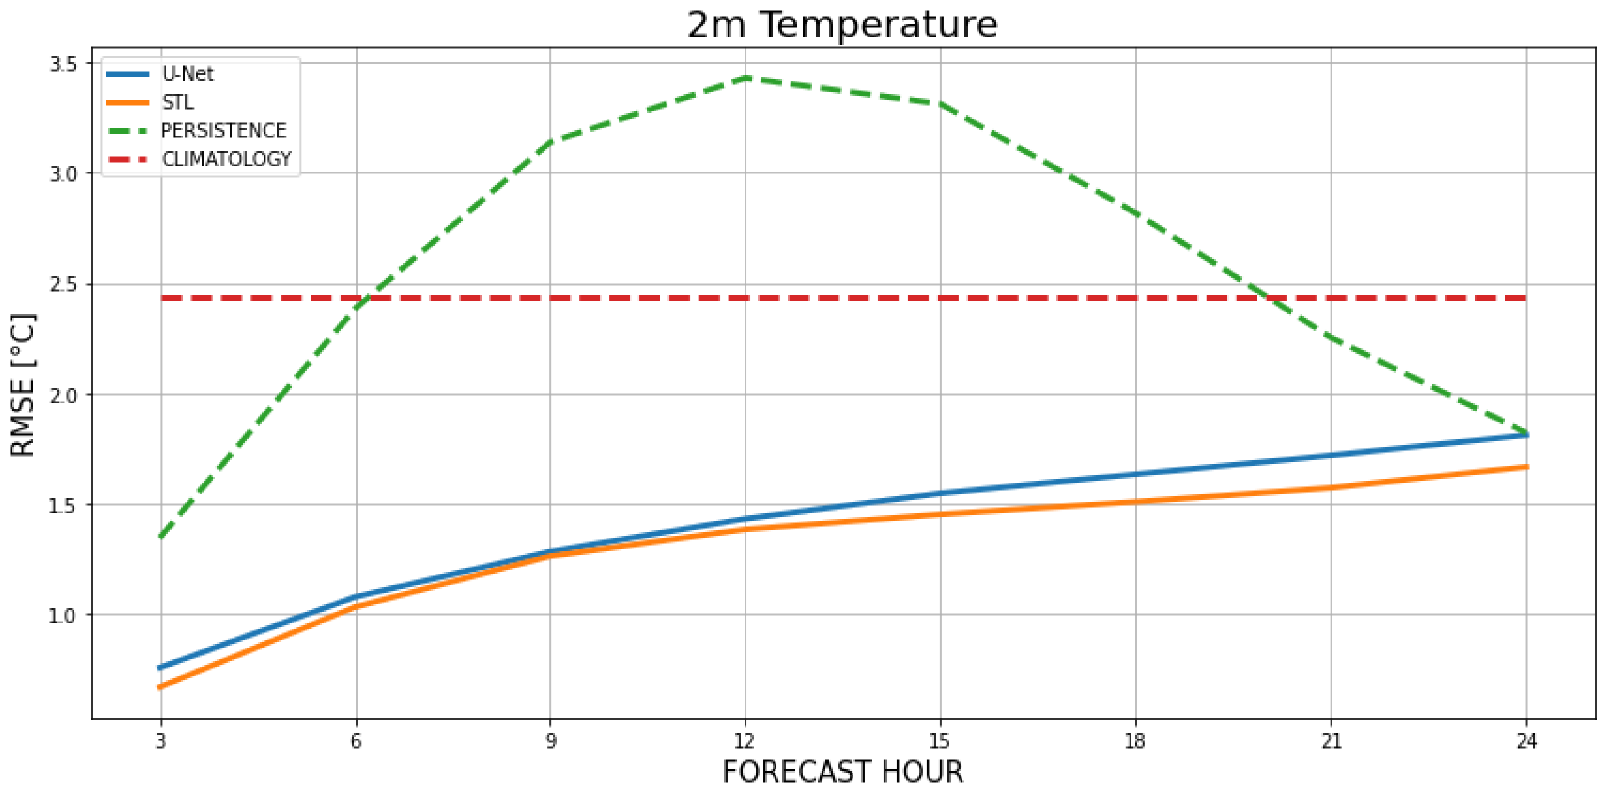

The first results seem very promising. Both models developed show metrics that are fully **comparable with those found in the literature and obtained by numerical methods**. In the last months of the project, the models will be optimised and compared with other ML models as well as with numerical models. In addition, thanks to Cineca's Mistral platform, the results obtained will be **compared with real data** from the observation stations.

WHAT'S NEXT?

From forecast to social impact

WEATHER4ENERGY

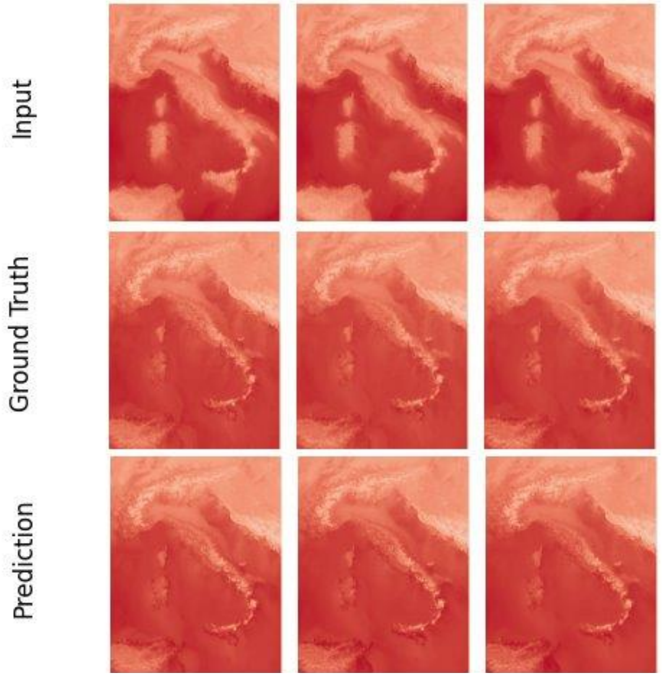

The project focuses on the efficient **translation of meteorological information into impacts**, providing accurate predictions of solar and wind power generation, power line capacity, hydrogeological hazards and potential threats to road infrastructure under future climate scenarios. The project is based on advanced meteorological models and involves the development of an innovative framework for translating meteorological information into concrete impacts, supported by high-performance infrastructure and cloud computing provided by the partnership with Cineca. This project is funded by the Innovation Grants within the scope of Centro Nazionale for a total of 24 months.

Questions
Doubts
Curiosities

GRAZIE

CINECA

m.angelinelli@cineca.it

Matteo Angelinelli