MAELSTROM

Empowering weather & climate forecast:

ML Apps & Datasets

ML Workflow Tools

Hardware Systems

Peter Dueben (ECMWF)

EuroHPC Joint Undertaking

"The MAELSTROM project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955513. The JU receives support from the European Union's Horizon 2020 research and innovation programme and United Kingdom, Germany, Italy, Luxembourg, Switzerland, Norway".

Co-ordinated by ECMWF

Why would machine learning help for weather and climate prediction?

Predictions of weather and climate are difficult, as the Earth system is huge, complex and chaotic, and as the resolution of our models is limited.

However, we have several hundred peta-bytes of Earth system data from observations and model output.

There are many application areas for machine learning in numerical weather predictions.

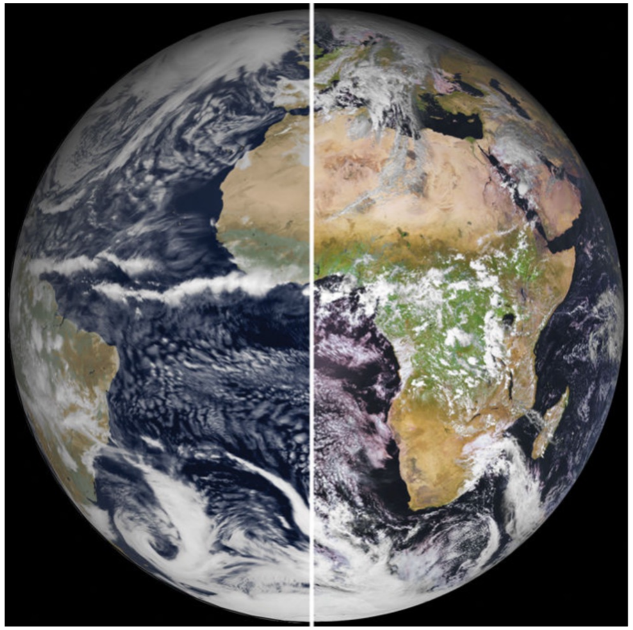

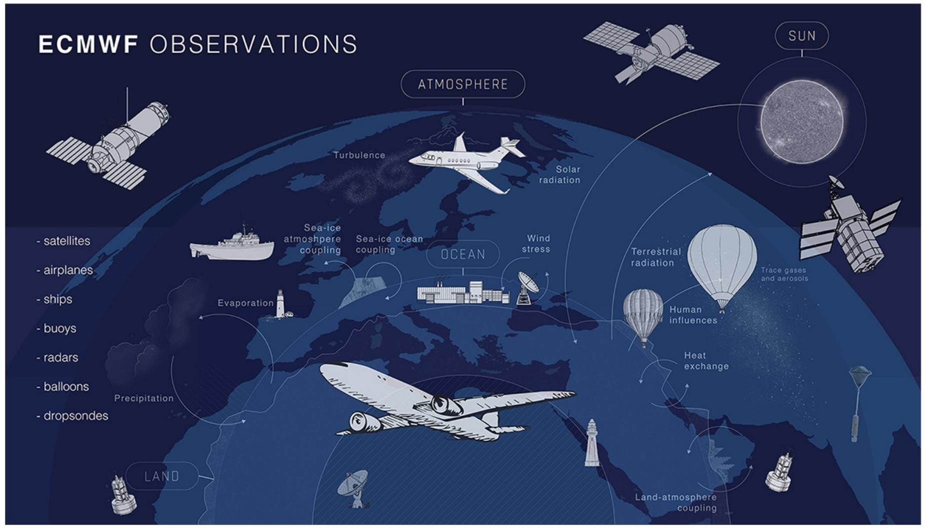

MAELSTROM – Why now?

Increase in data volume

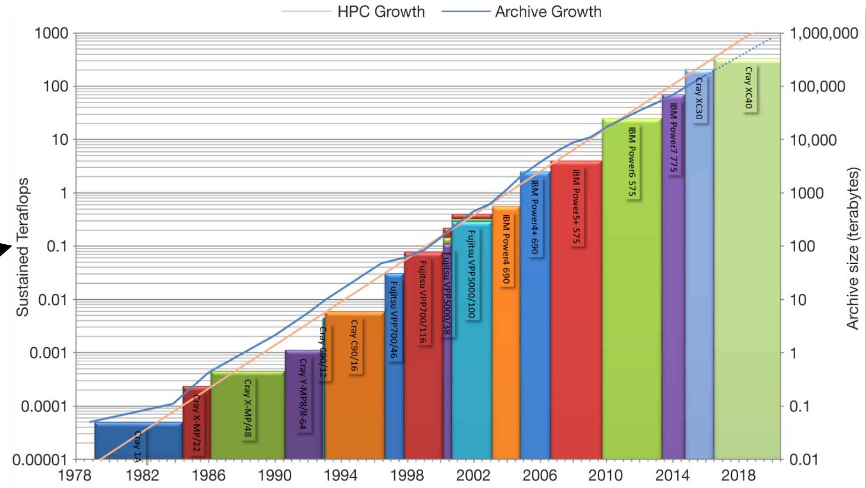

New computing hardware

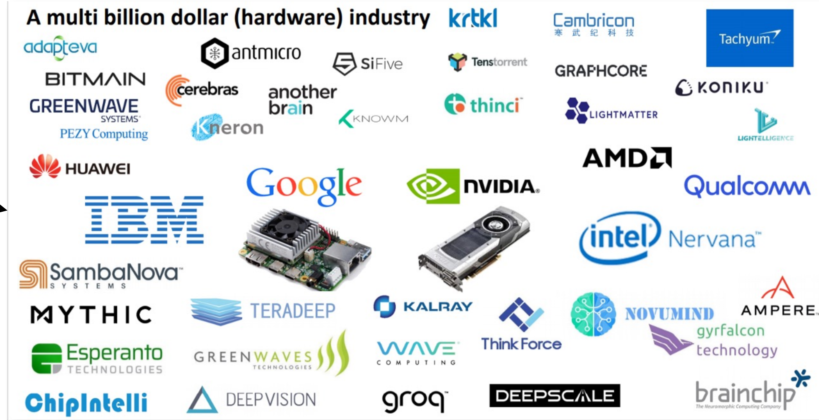

New machine learning software

TensorFlow

Keras

PyTorch

Increase in knowledge

A myriad of options…

A myriad of options for machine learning approaches

Dense Neural Networks, LSTMs, ConvGru, Attention Layers, Transformer networks, # of hidden layers, different normalisation of inputs, batch normalisation, tanh, relu, gelu, softplus, elu, selu, leaky relu, softmax, sigmoid function, generative adversarial networks, recurrent neural networks, encoding/decoding networks, random forests, boosting methods, clustering techniques, singular vector decomposition, causal discovery, ablation studies, root mean square error, variational auto encoder, gradient descent, stochastic gradient decent, adagrad, adadelta, RMSprop, Adam, # of epochs, # of batches, learning rate, overfitting, dropout, Bayesian networks, Gaussian processes, half precision, sparse networks….

+ a myriad of options for machine learning hardware

CPUs, RISC-V, GPUs from different vendors, Tensorcores, TPUs, FPGAs, ASICs, European Processor Initiative, GRAPHCORE, Sambanova, CERVEST, double precision, single precision, half precision, Bfloat16, Bfloat32, Cloudcomputing…

= confused scientists

The MAELSTROM Objectives

O1: To open W&C predictions as a new usage domain for machine learning applications that can exploit exaflop performance.

O2: To develop the optimal software environment to develop exascale-ready machine learning tools that can be used across the workflow of W&C predictions.

O3: To optimise compute system designs for machine learning applications for W&C predictions at the node and system level and to transfer this knowledge to other machine learning applications that will use future EuroHPC systems.

The MAELSTROM Objectives

Specific Objective 1 (SO1): MAELSTROM will develop benchmark datasets for six selected ML applications that cover the entire workflow of W&C predictions.

Specific Objective 2 (SO2): MAELSTROM will develop production-ready machine learning solutions that are optimised for efficiency, scalability, and quality.

Specific Objective 3 (SO3): MAELSTROM will develop bespoke machine learning workflow tools for W&C applications that optimise collaborations between W&C, machine learning and HPC experts and allow for a prompt uptake and operational implementation of machine learning within W&C models as well as the performance benchmarking of machine learning solutions based on Deep500.

Specific Objective 4 (SO4): MAELSTROM will develop bespoke system-level architecture blueprints for ML.

The MAELSTROM approach
A co-design cycle

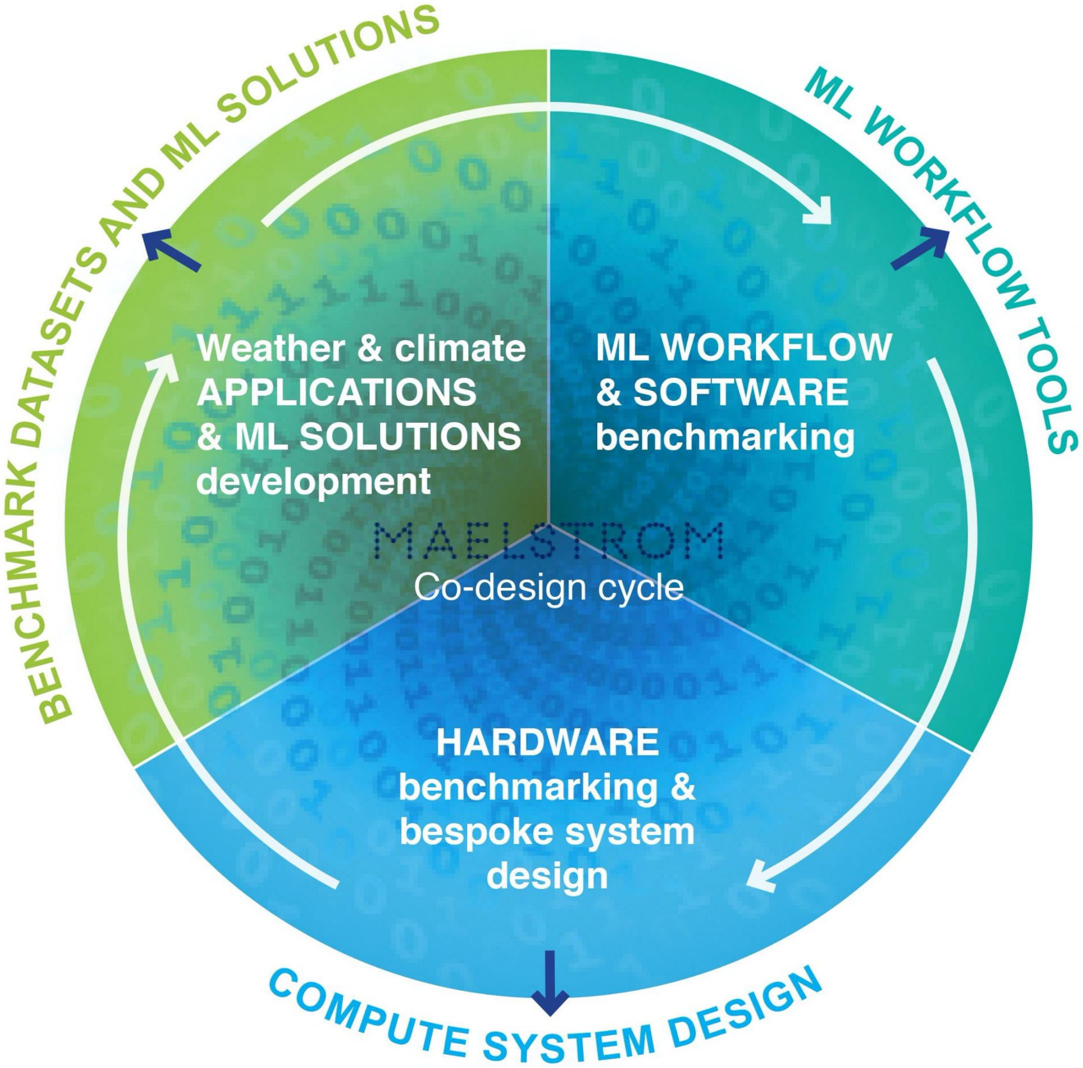

We are MAELSTROM

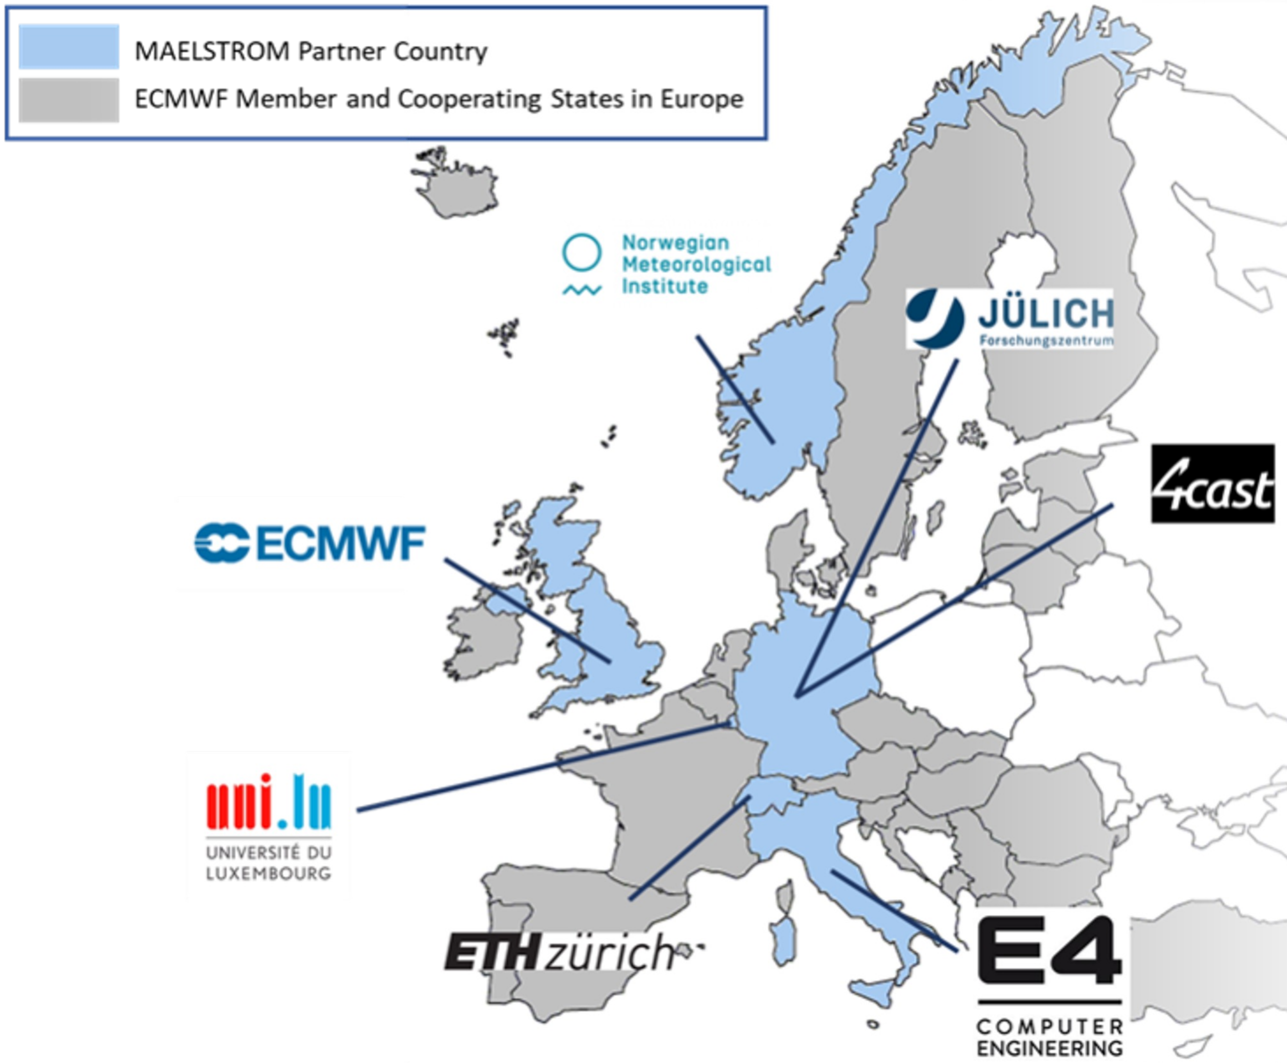

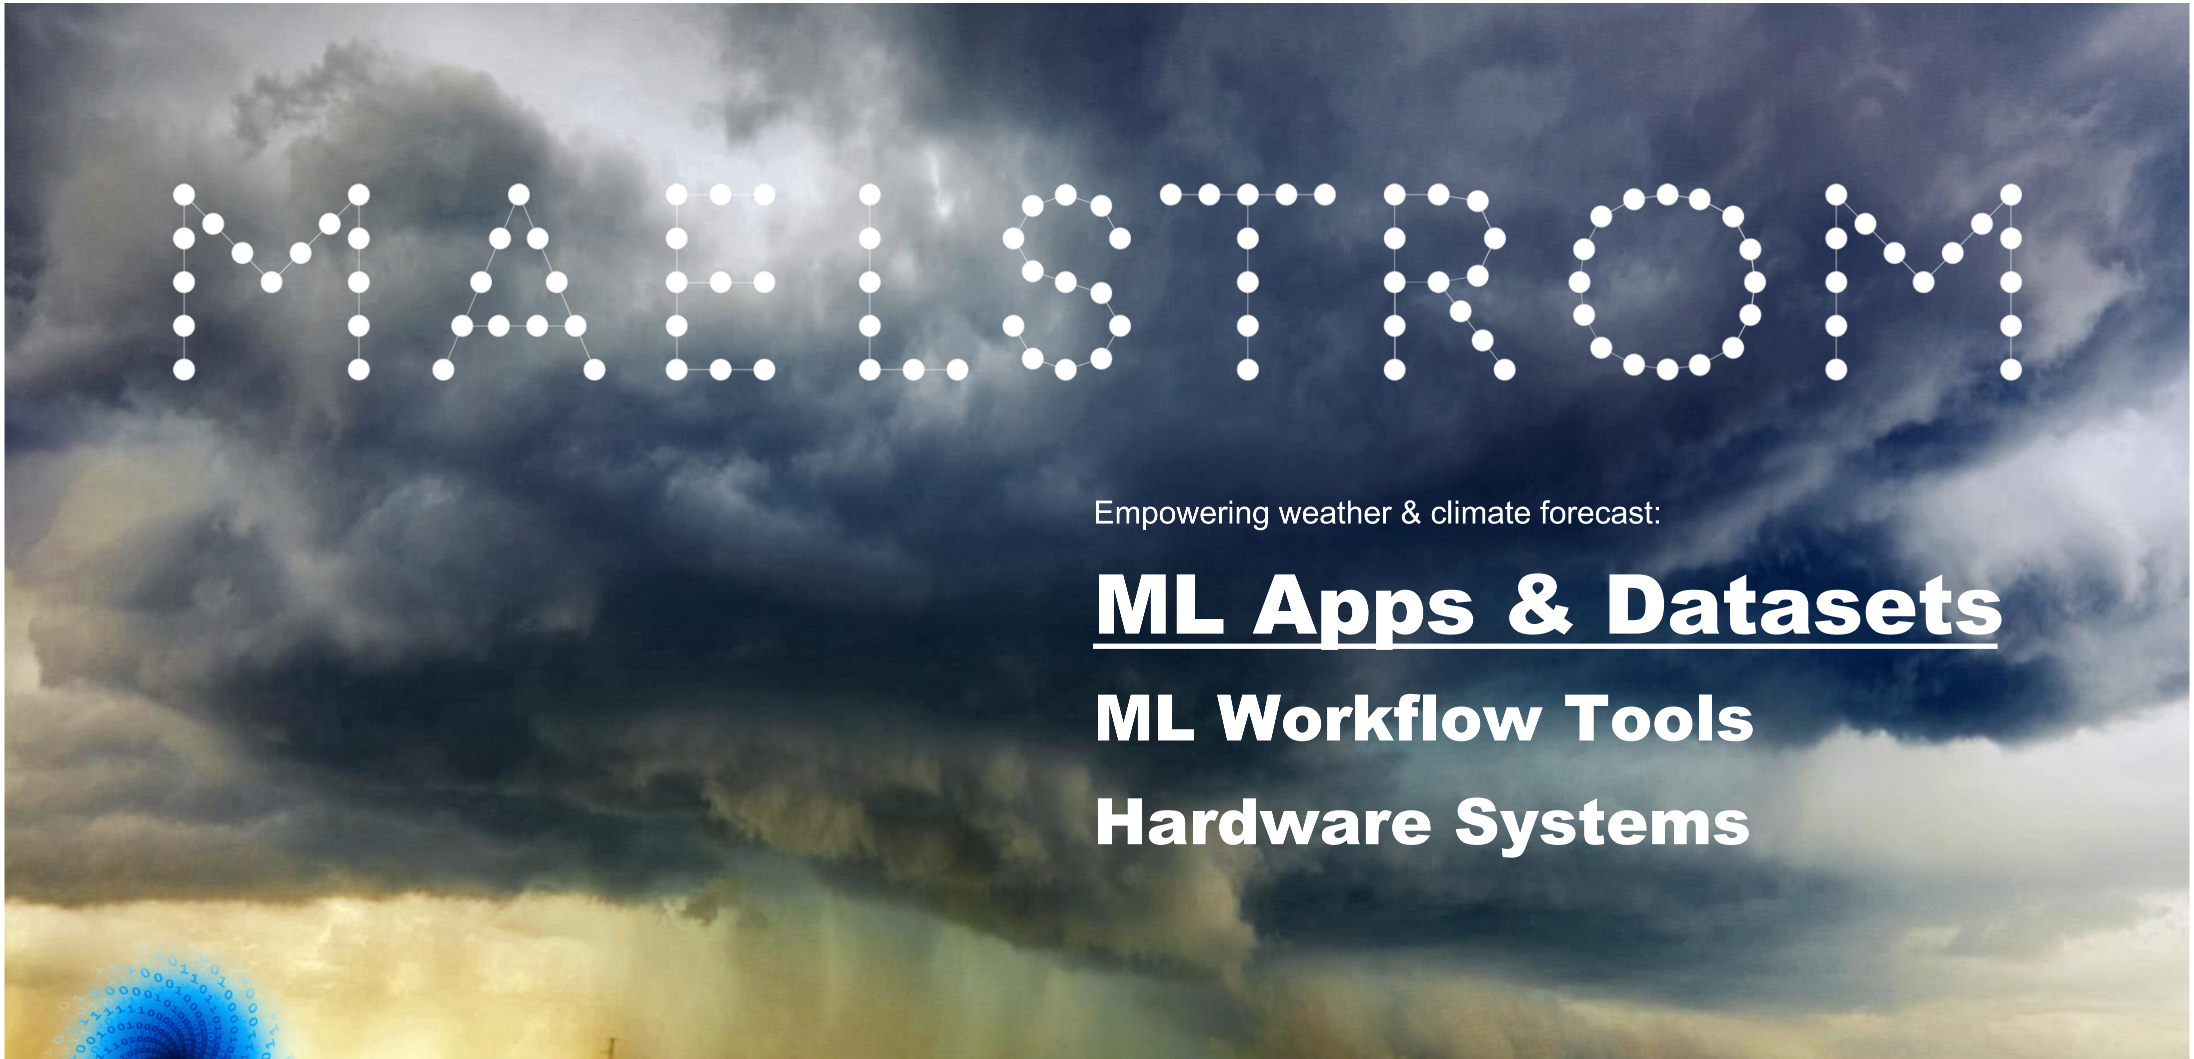

EuroHPC
Joint Undertaking

"The MAELSTROM project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955513. The JU receives support from the European Union's Horizon 2020 research and innovation programme and United Kingdom, Germany, Italy, Luxembourg, Switzerland, Norway".

Co-ordinated by
ECMWF

Motivation: ML for weather forecasts?

A1: Use citizen observations

A4: Post-process ensemble simulations

A6: Improve weather predictions at windfarms

A2: Use tweets as weather sensors

A3: Emulate model components

A5: Down-scale precipitation forecasts

Observations
$O(10^8)$ data points → Data-assimilation
$O(10^7)$ data points → Numerical weather forecasts
$O(10^{10})$ data points → Post-processing of model output
$O(10^{10})$ data points → Product generation
$O(10^5)$ data points

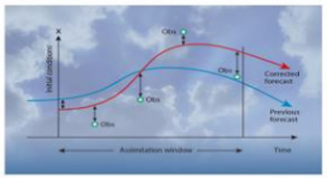

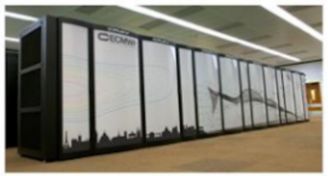

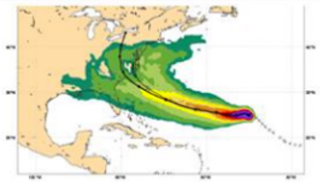

Objective: open W&C prediction as a new domain for ML applications that exploit exaflop performance

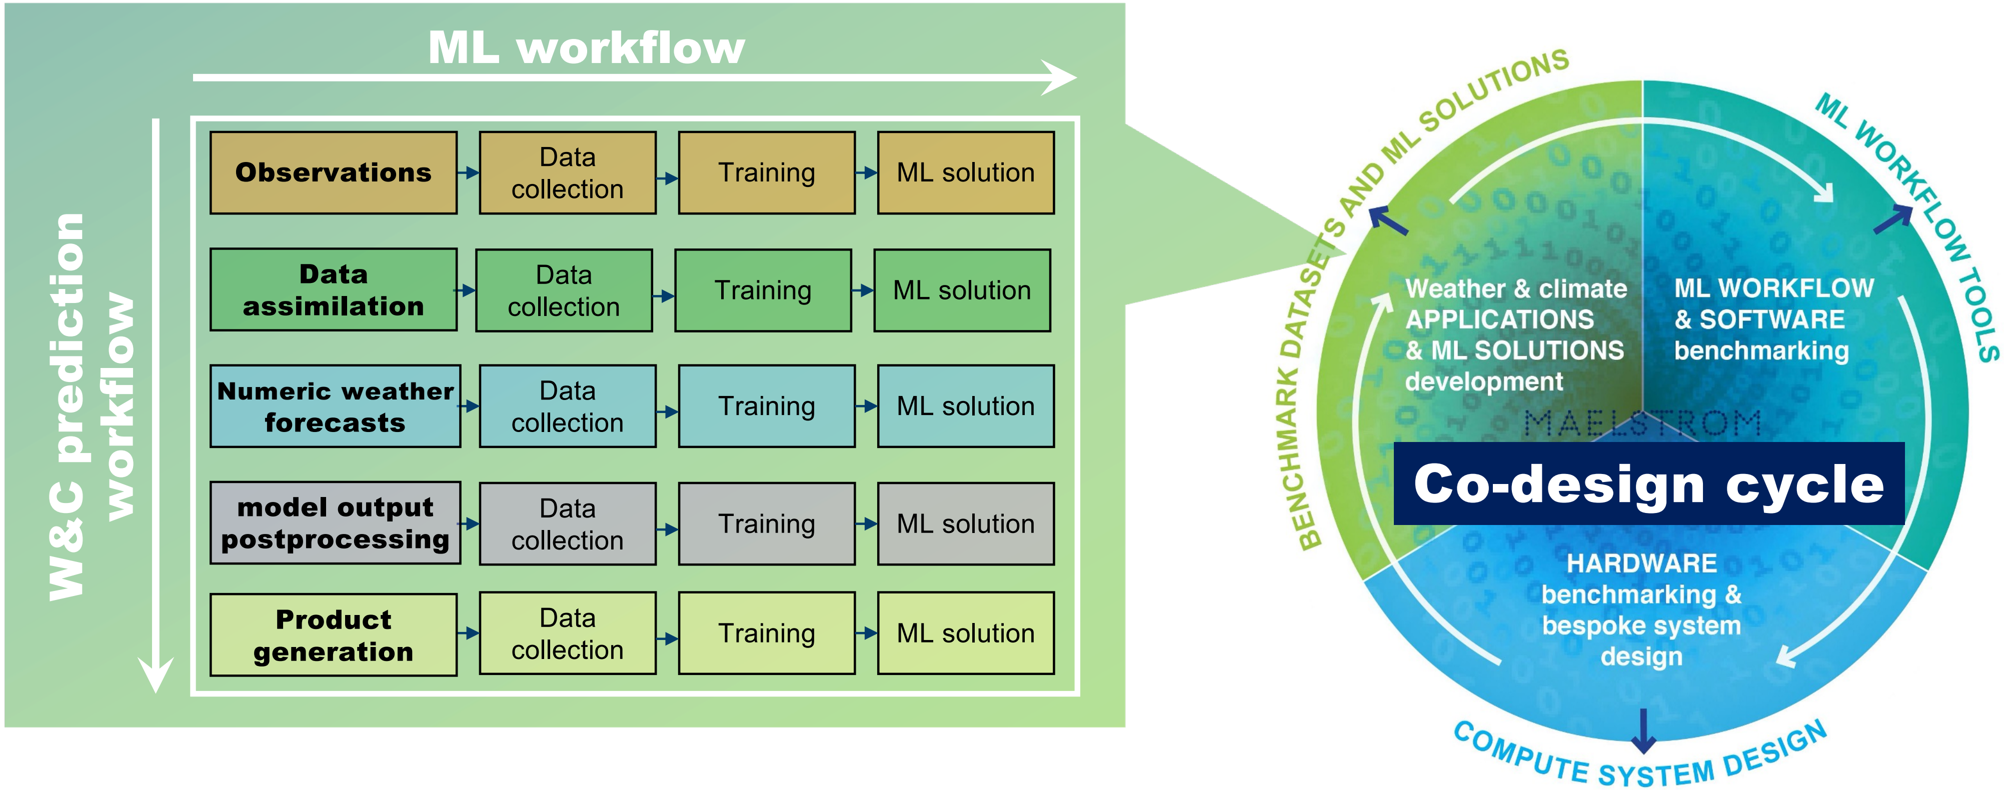

Make developments comparable via benchmark datasets

Benchmark datasets include:

- A problem statement
- Data that is available online
- Python code or Jupyter notebooks
- A reference machine learning solution
- Quantitative evaluation metrics
- Visualisation, diagnostics and robustness tests
- Computational benchmarks

Benchmark datasets are useful because:

- They allow a quantitative evaluation of machine learning approaches
- They reduce data access and help scientists to get access to relevant data
- They allow for a separation of concerns between domain sciences and machine learning experts
- They allow for a separation of concerns between domain sciences and HPC experts

Missing machine learning benchmark datasets for atmospheric sciences

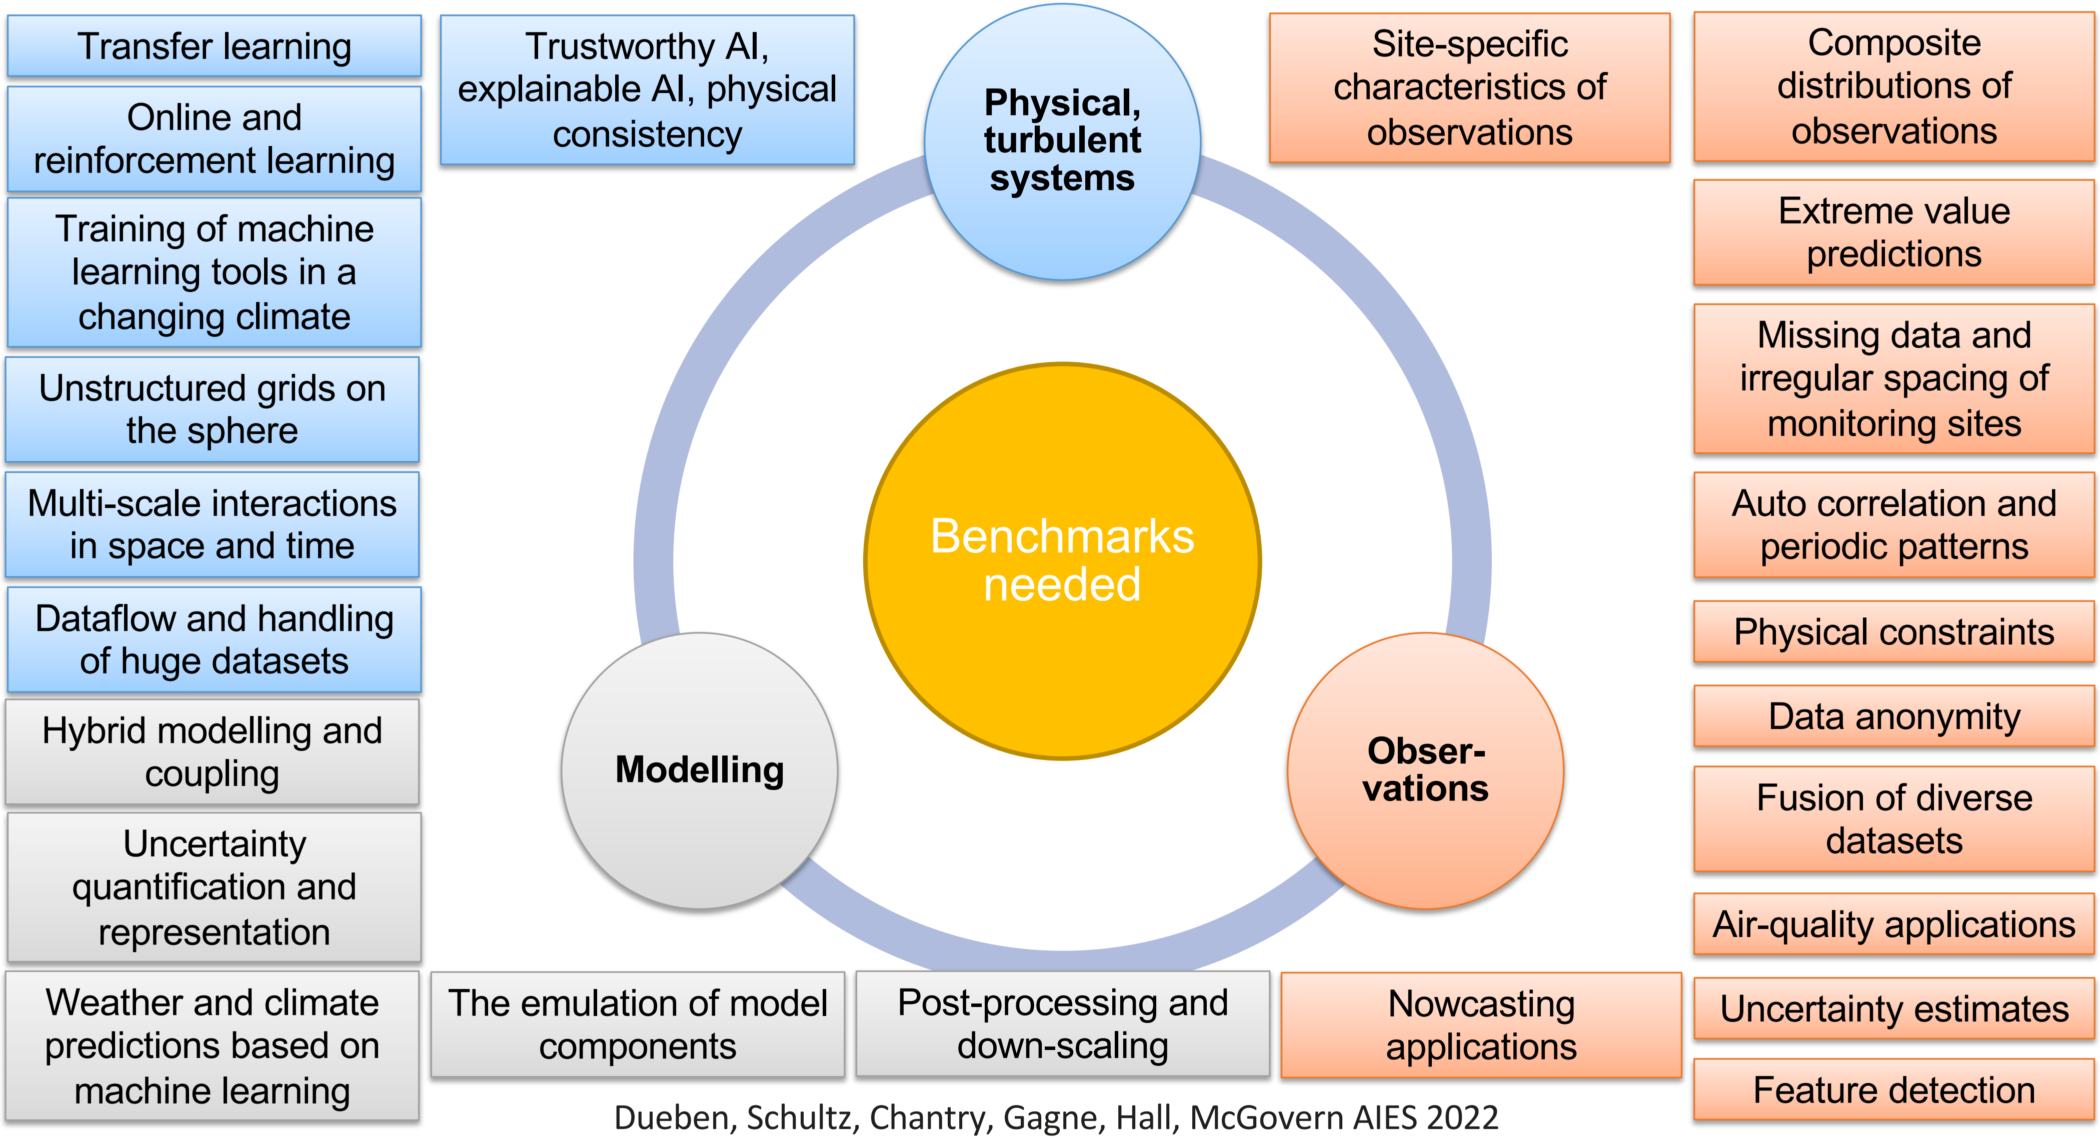

Dueben, Schultz, Chantry, Gagne, Hall, McGovern AIES 2022

Develop machine learning benchmark datasets

The infrastructure based on S3 buckets for data storage were set up on the ECMWF data cloud.

CliMetLab plugins are used to manage the downloading of the dataset. The plugins have been created for the six applications.

Jupyter notebooks have been created to explore the datasets and demonstrate simple machine learning solutions to act as first benchmarks

Hardware performance testing

```
!pip install climetlab climetlab-maelstrom-radiation
import climetlab as cml
cmlds = cml.load_dataset('maelstrom-radiation')
ds = cmlds.to_xarray()
```

Application	Pip package name	CML dataset name
A1: Postprocessing	climetlab-maelstrom-yr	'maelstrom-yr'
A3: Radiation	climetlab-maelstrom-radiation	'maelstrom-radiation'
A4: ENS10	climetlab-maelstrom-ens10	'maelstrom-ens10'
A5: Downscaling	climetlab-maelstrom-downscaling	'maelstrom-downscaling'
A6: Power production	climetlab-maelstrom-power-production	'maelstrom-constants-a-b' 'maelstrom-power-production' 'maelstrom-weather-model-level' 'maelstrom-weather-pressure-level' 'maelstrom-weather-surface-level'

Perform tests to trade efficiency, quality and speed

Tests with reduced numerical precision

Tests with different machine learning software libraries

Scalability tests

Tests with different machine learning architectures

A long while after 2018…
The perspective of full machine learning models for weather and climate

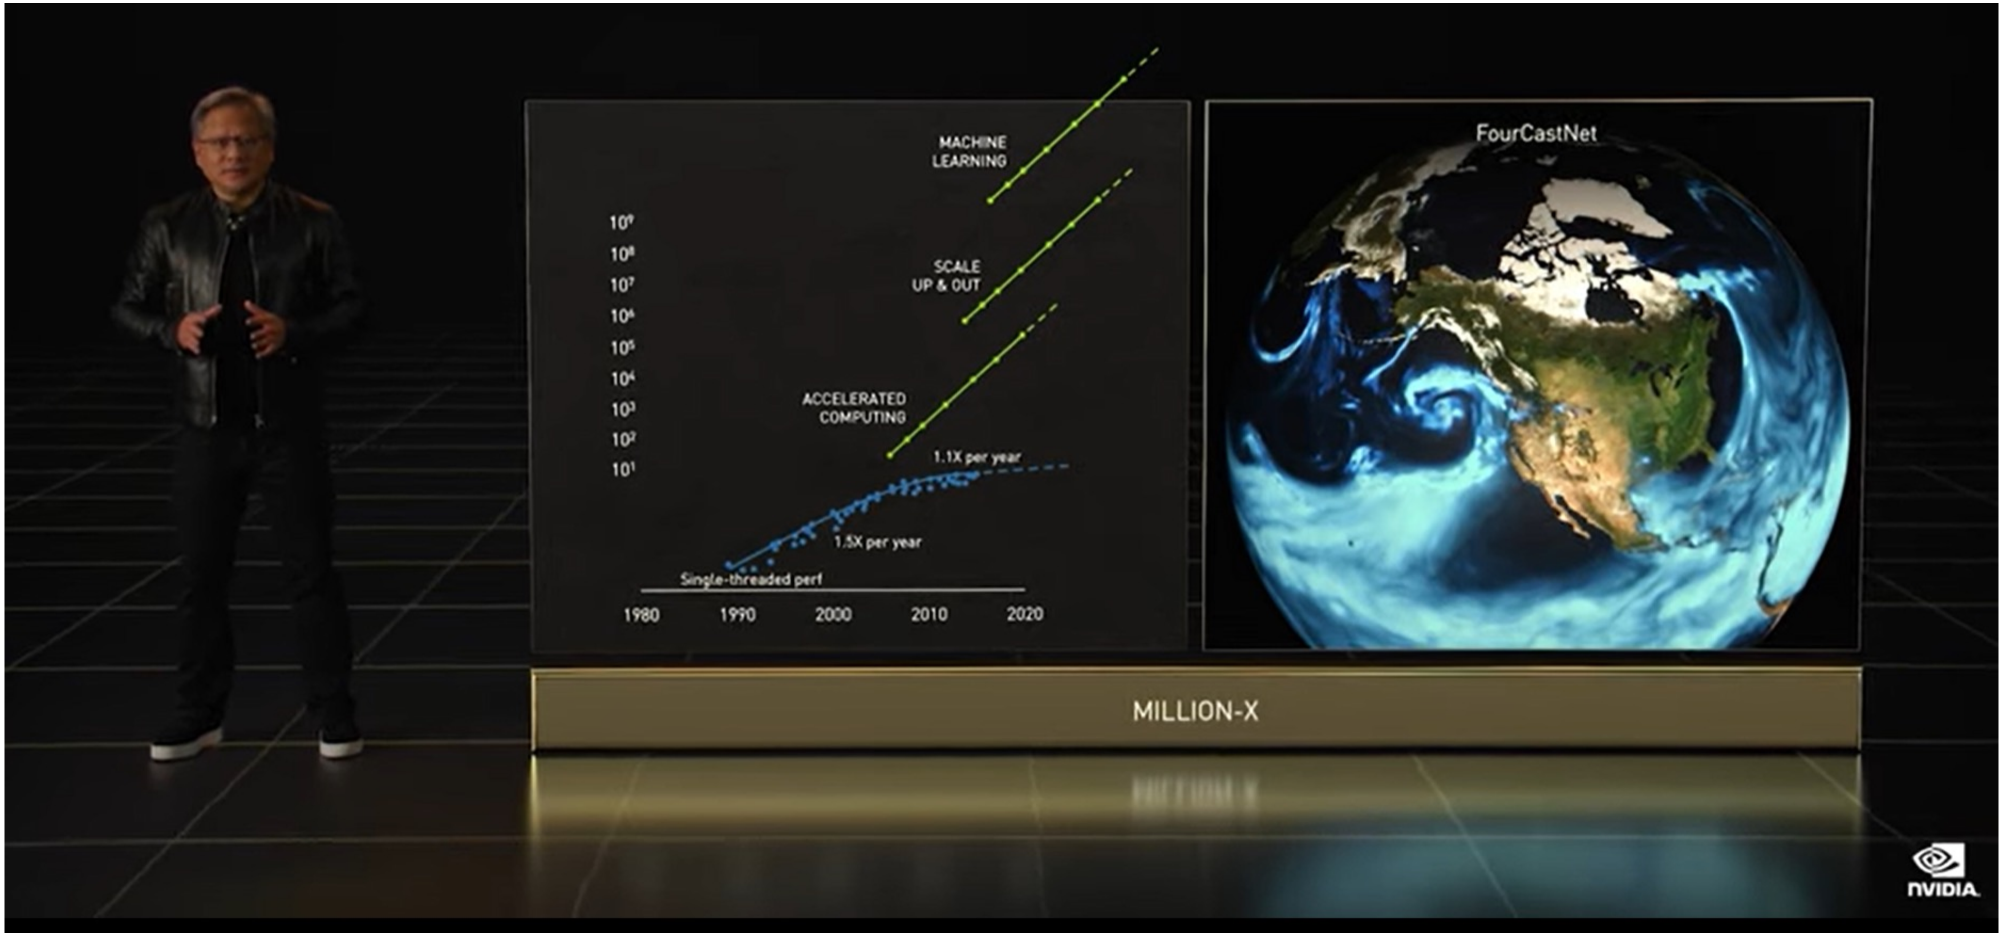

NIVIDA's Earth-2 is coming with FourCastNet

2022-today: The machine learning revolution

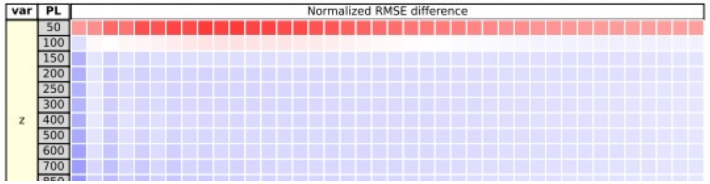

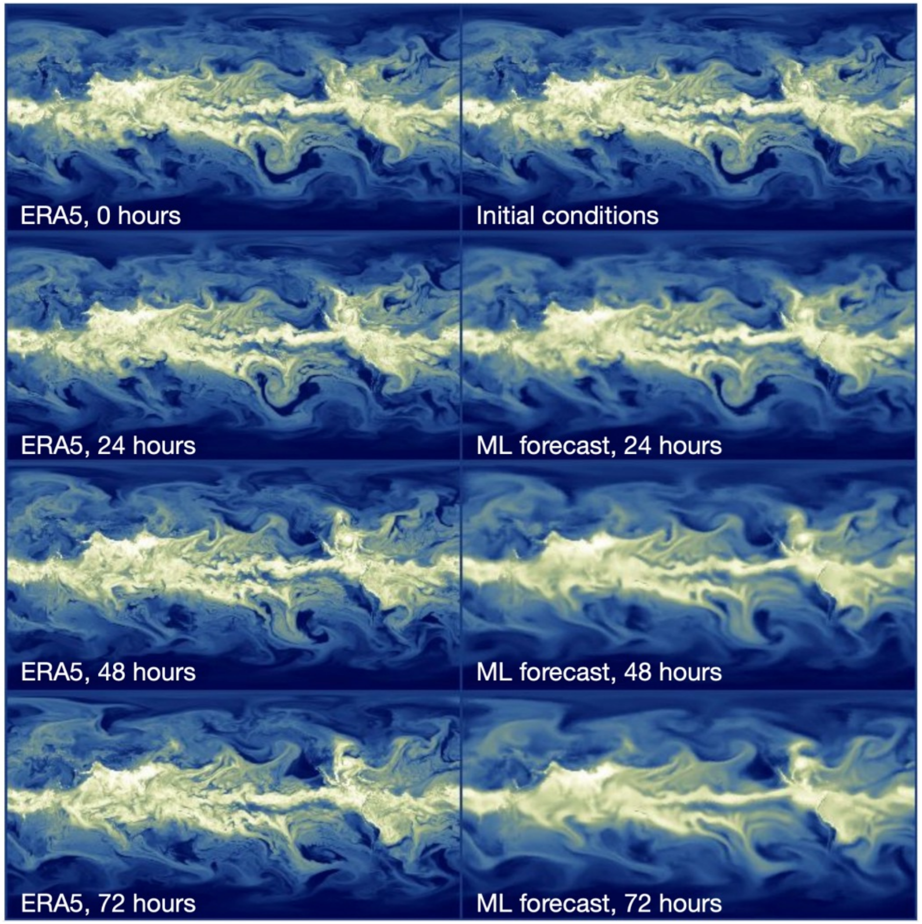

GraphCast from Google/Deepmind and Fourcastnet from NVIDIA are beating conventional weather forecast model in deterministic scores and are orders of magnitudes faster.

But how do these models actually work?

They get the best results when using very large timesteps.

They are trained for a small Root Mean Square Error.
→ They smear out for large lead times.

Many questions remain:
Can the models extrapolate?
Can they represent extreme events?
Can they learn uncertainty?
Can they be trained from observations?
Can they represent physical consistency?

Images from Keisler (2022)

2022-today: The machine learning revolution

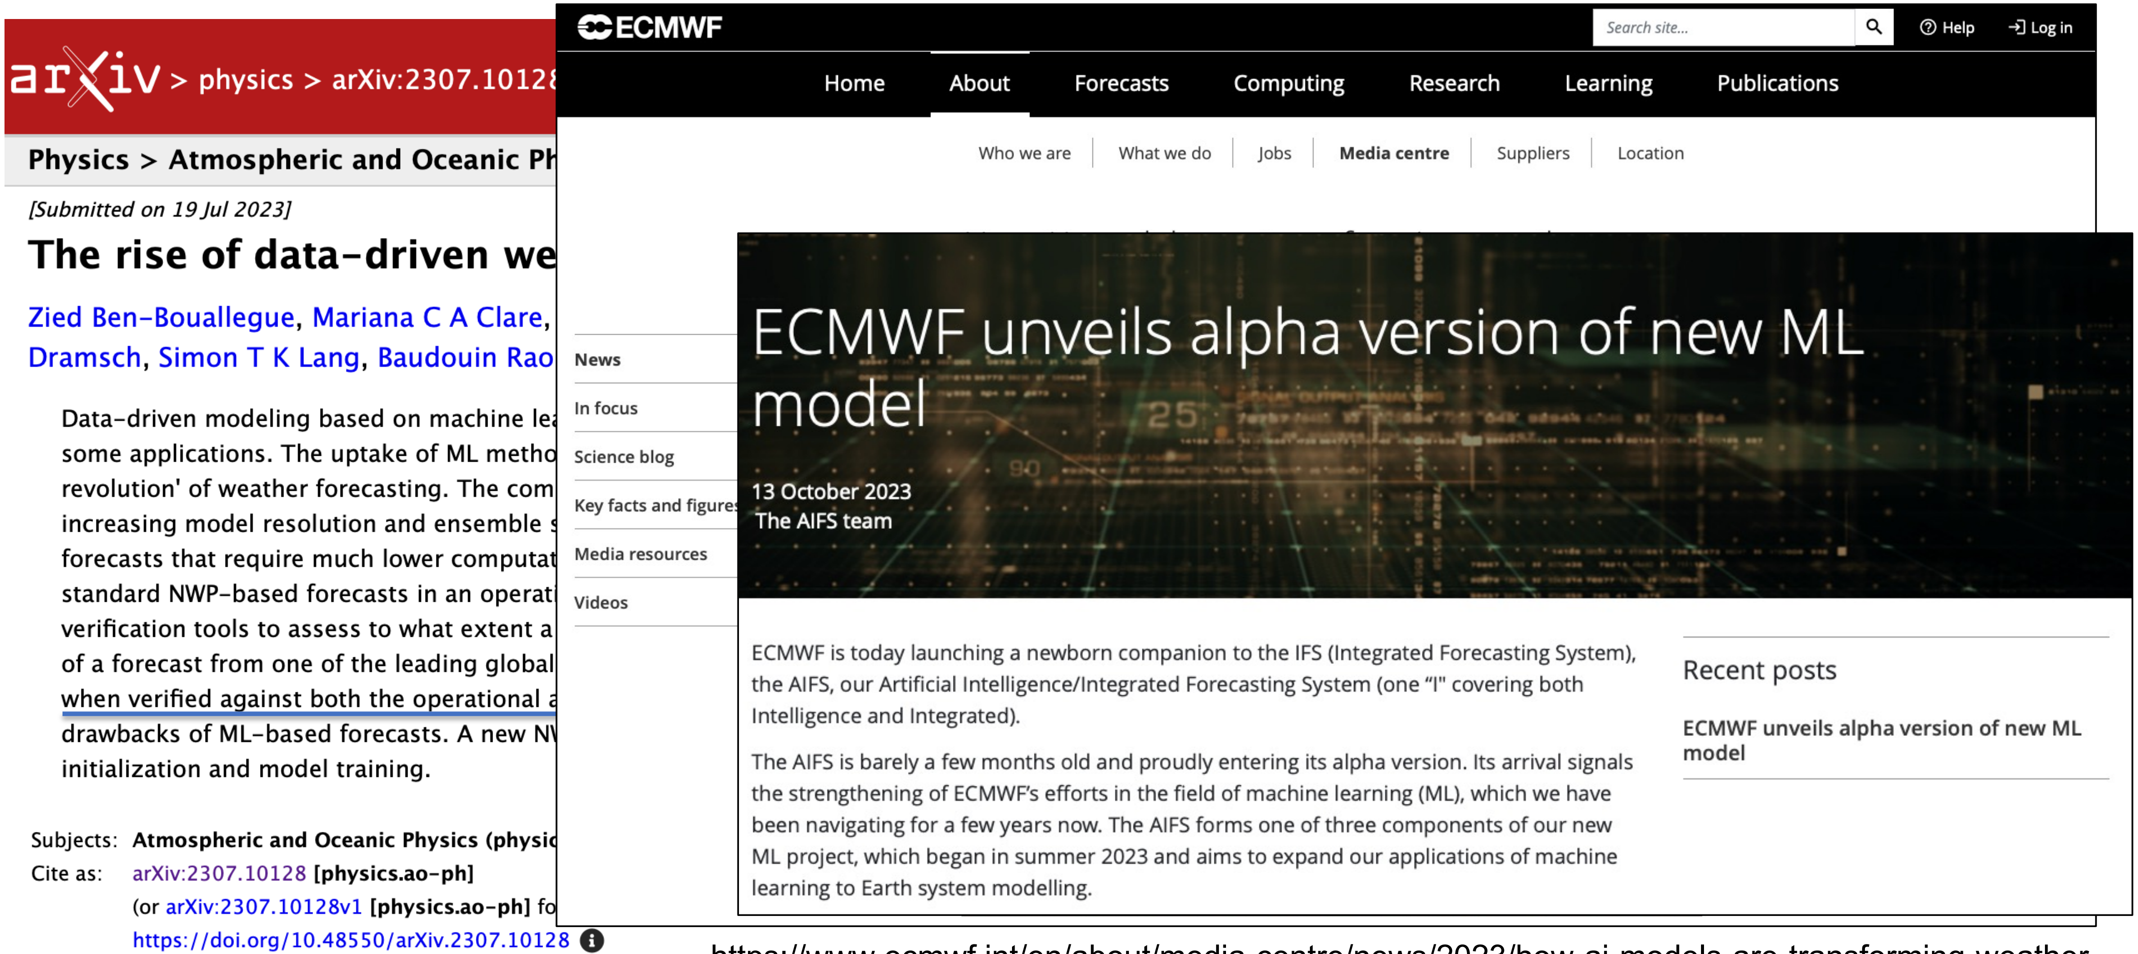

arXiv > physics > arXiv:2307.1012

Physics > Atmospheric and Oceanic Ph

[Submitted on 19 Jul 2023]

The rise of data-driven we

Zied Ben-Bouallegue, Mariana C A Clare, Dramsch, Simon T K Lang, Baudouin Rao

Data-driven modeling based on machine le some applications. The uptake of ML metho revolution' of weather forecasting. The com increasing model resolution and ensemble forecasts that require much lower computat standard NWP-based forecasts in an operat verification tools to assess to what extent a of a forecast from one of the leading global when verified against both the operational drawbacks of ML-based forecasts. A new N initialization and model training.

Subjects: **Atmospheric and Oceanic Physics (physi**
Cite as: arXiv:2307.10128 **[physics.ao-ph]**
(or arXiv:2307.10128v1 **[physics.ao-ph]** fo
https://doi.org/10.48550/arXiv.2307.10128

Submission history
From: Zied Ben Bouallegue [view email]
[v1] Wed, 19 Jul 2023 16:51:08 UTC (18,531 KB)

ECMWF
Search site... Help Log in
Home About Forecasts Computing Research Learning Publications
Who we are | What we do | Jobs | **Media centre** | Suppliers | Location

News
In focus
Science blog
Key facts and figure
Media resources
Videos

ECMWF unveils alpha version of new ML model

13 October 2023
The AIFS team

ECMWF is today launching a newborn companion to the IFS (Integrated Forecasting System), the AIFS, our Artificial Intelligence/Integrated Forecasting System (one "I" covering both Intelligence and Integrated).

The AIFS is barely a few months old and proudly entering its alpha version. Its arrival signals the strengthening of ECMWF's efforts in the field of machine learning (ML), which we have been navigating for a few years now. The AIFS forms one of three components of our new ML project, which began in summer 2023 and aims to expand our applications of machine learning to Earth system modelling.

Recent posts

ECMWF unveils alpha version of new ML model

https://www.ecmwf.int/en/about/media-centre/news/2023/how-ai-models-are-transforming-weather-forecasting-showcase-data
https://www.ecmwf.int/en/about/media-centre/aifs-blog/2023/ECMWF-unveils-alpha-version-of-new-ML-model

Representation learning and a Machine Learned Foundation Model next?

AtmoRep: A stochastic model of atmosphere dynamics using large scale representation learning

Christian Lessig[1*], Ilaria Luise[2], Bing Gong[3], Michael Langguth[3], Scarlet Stadtler[3], Martin Schultz[3]

[1]Department of Computer Science, Otto-von-Guericke-Universität Magdeburg, Universitätsplatz 2, Magdeburg, Germany.
[*]now at the European Centre for Medium Range Weather Forecasting, Robert-Schumann-Platz, Bonn, Germany.
[2]CERN, European Center for Nuclear Research, Esplanade des Particules 1, Meyrin, Switzerland.
[3]Jülich Supercomputing Centre, Forschungszentrum Jülich, Wilhelm-Johnen-Str., Jülich, Germany.

Contributing authors: christian.lessig@ovgu.de; ilaria.luise@cern.ch; b.gong@fz-juelich.de; m.langguth@fz-juelich.de; s.stadtler@fz-juelich.de; m.schultz@fz-juelich.de;

Abstract

The atmosphere affects humans in a multitude of ways, from loss of life due to adverse weather effects to long-term social and economic impacts on societies. Computer simulations of atmospheric dynamics are, therefore, of great importance for the well-being of our and future generations [1, 2]. Classical numerical models of the atmosphere, however, exhibit biases due to incomplete process descriptions and they are computationally highly demanding [1]. Very recent AI-based weather forecasting models [3–7] reduce the computational costs but they lack the versatility of conventional models and do not provide probabilistic predictions. Here, we propose AtmoRep, a novel, task-independent stochastic computer model of atmospheric dynamics that can provide skillful results for a wide range of applications. AtmoRep uses large-scale representation learning from artificial intelligence [8, 9] to determine a general description of the highly complex, stochastic dynamics of the atmosphere from the best available estimate of the system's historical trajectory as constrained by observations [10]. This is enabled by a novel self-supervised learning objective and a unique ensemble that samples from the stochastic model with a variability informed by the one in

EuroHPC
Joint Undertaking

"The MAELSTROM project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955513. The JU receives support from the European Union's Horizon 2020 research and innovation programme and United Kingdom, Germany, Italy, Luxembourg, Switzerland, Norway".

Co-ordinated by
ECMWF

Goal

Workflow tools for ML in W&C used for development and implementation

Software for ML development on HPC

mantik mlflow CliMetLab DEEP500 aCeML NoPFS Near-optimal Pre-Fetching System JUBE BENCHMARKING ENVIRONMENT

- Promote collaboration
- Ease the ML workflow
- Unified access to hardware systems

Requirements of Mantik

mantik

Reproducibility of ML solutions

- ✓ Recording of model input parameters, metrics
- ✓ Saving and loading of trained models

Interface to Compute Resources

- ✓ Abstract away infrastructure
- ✓ Unified access to compute resources (HPC)
- ✓ Run jobs from Platform
- SOON Inference of trained models

Sharing and Recommending

- ✓ Sharing ML solutions
- ✓ Work in common project

Real-time tracking with mlflow

Secured access to HPC via the GUI through cluster interfaces:

UNICORE

FIRECREST

- ➢ Labels available for ***Projects***, ***Code*** and ***Experiments***
- ➢ Users can form ***Groups*** and ***Organisations***

4cast

The Mantik Web Interface

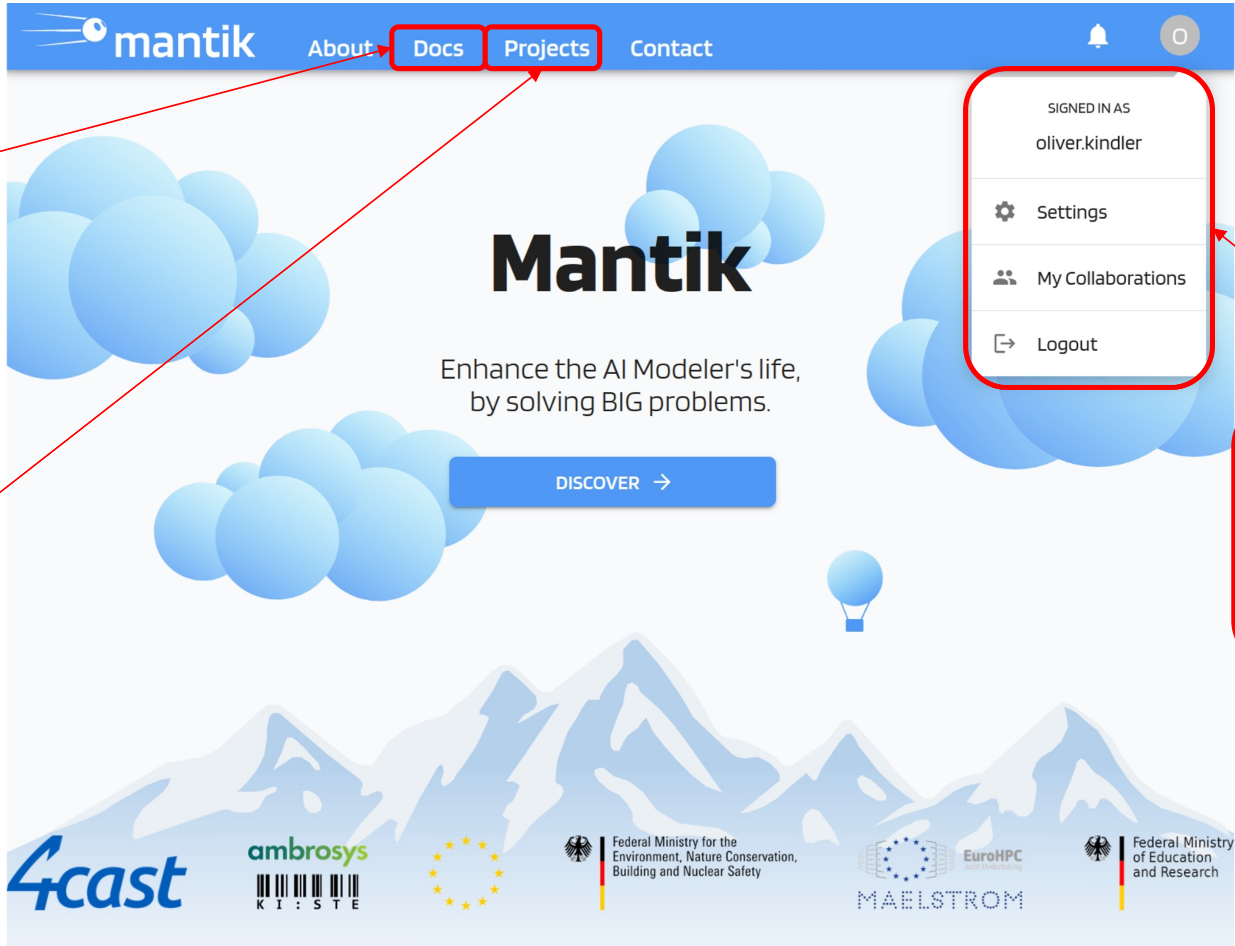

Quickstart Guides

Project Tutorials

CLI Documentation

Set up HPC connection

Manage your account

Form Groups

See all available *Projects*

To get started, visit:
https://cloud.mantik.ai

Projects Page

Search *Projects* by *Labels*

Scopes of Labels are predefined

Labels are grouped in different classes

New Labels can be suggested via a Service Desk

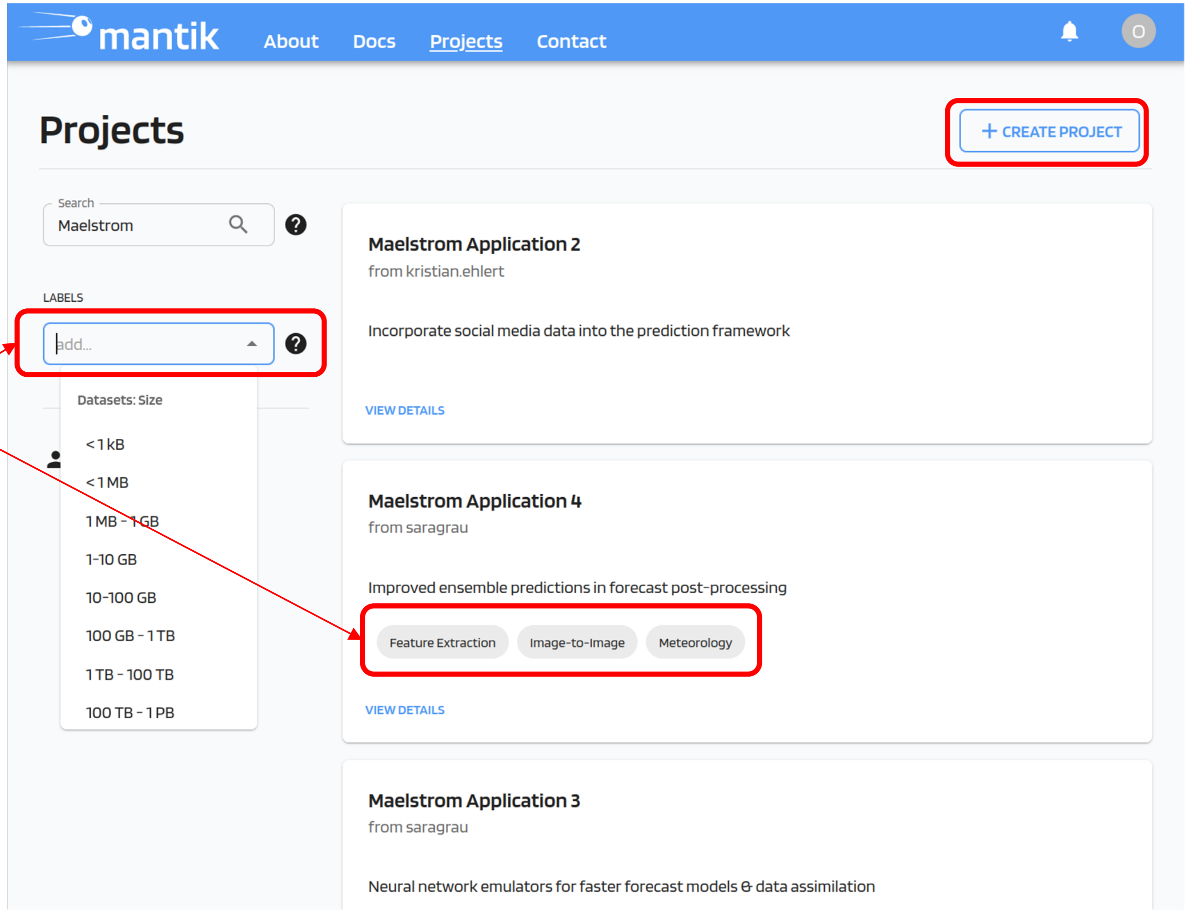

To get started, visit: https://cloud.mantik.ai

Submitting a *Run*

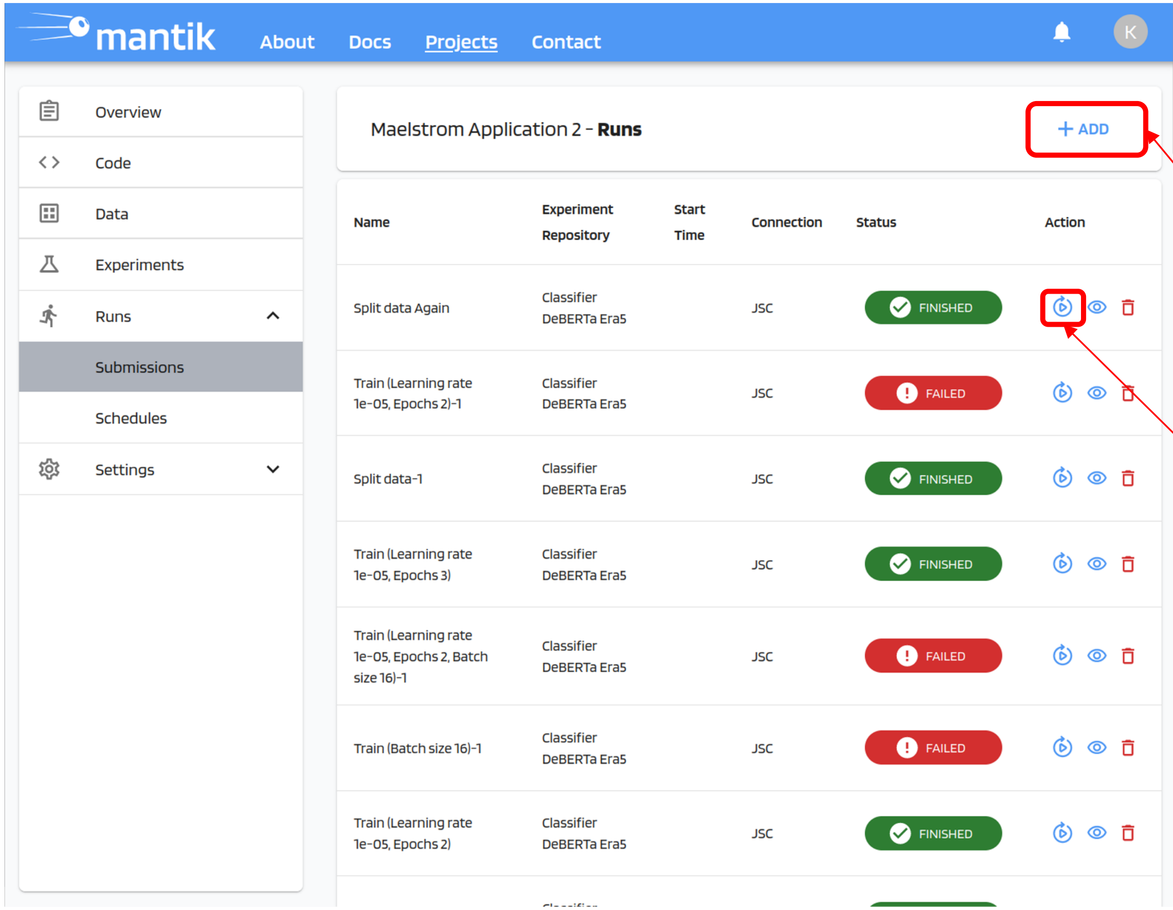

Every Submission of a compute job to a HPC cluster is called a *Run*

Runs are grouped in *Experiments (later more)*

Re-Run Button

Change parameters and repeat an earlier *Run*

To get started, visit: https://cloud.mantik.ai

EuroHPC Joint Undertaking

"The MAELSTROM project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955513. The JU receives support from the European Union's Horizon 2020 research and innovation programme and United Kingdom, Germany, Italy, Luxembourg, Switzerland, Norway".

Co-ordinated by ECMWF

WP3: Hardware benchmarking and bespoke system design

Objectives:

- perform system-level benchmarking for computation and data management of ML solutions
- develop customised reference system designs for ML solutions
- optimise compute system designs for ML applications for W&C predictions

How?

- enabling developers of ML solutions and workflow tools to access different relevant technologies to test their tools within the MAELSTROM co-design framework
- providing a continuous and reliable feedback on architectural features
- facilitating discussion between application developers and hardware experts for informed design choices
- Using benchmarking data to enhance knowledge and understanding of ML applications performance

Why?

- For example, ECMWF will add a ML benchmark to the next procurement tests
- EuroHPC systems see heavy ML workloads on the GPU partitions

Computing Systems used for last benchmarking (E4 Systems)

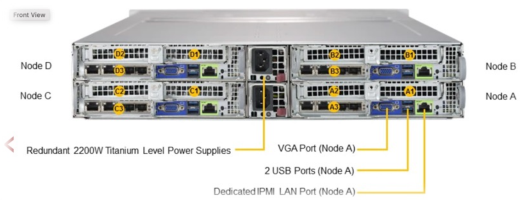

E4 Intel Cluster

1 login node

4 nodes in a single chassis

- OS RHEL release 8.6
- Dual Socket with 2x Intel Xeon Gold 6226R @ 2.90 GHz, 16-Core Processor
- 192 GB RAM

2 nodes single server

- OS RHEL release 8.5
- Dual socket with 2x Intel Xeon Gold 6326R, 16-Core Processor
- 512 GB RAM
- 1x NVIDIA A100 GPU (per node)

3 TB NVMe for each node

Infiniband 100 Gb/s Network

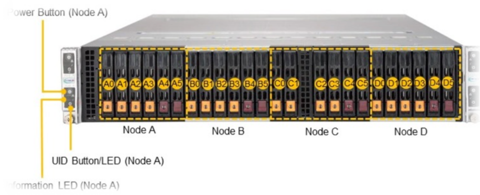

E4 AMD Cluster

1 login node

4 nodes in a single chassis

- OS RHEL release 8.6
- Dual socket with 2x AMD EPYC “Milan” 7453, 32-Core Processor
- 256 GB RAM

2 nodes single server

- OS RHEL release 8.5
- Dual socket with 2x AMD EPYC 7313, 16-Core Processor
- 512 GB RAM
- 1x AMD Instinct Mi100 GPU (per node)

Infiniband 100 Gb/s Network

Computing Systems used for last benchmarking (JSC Systems)

- **JUWELS Cluster (2018)**
  - 2511 computing nodes (2 x Skylake)
  - 48 GPU nodes (4 x NVIDIA V100 with NVLINK)
  - Mellanox 100 Gbit/s Fat Tree Topology (1:2 blocking factor)
  - 12 PFLOPs

- **JUWELS Booster (2020)**
  - 936 compute nodes,
  - Each: 4 NVIDIA A100 GPUs, 4 HDR-200 adapters
  - Mellanox HDR-200 InfiniBand in DragonFly+ topology
  - 73 PFLOPs

Top500 Nov-2020:
- #1 Europe
- #7 World
- #1* Green500

Outcomes of the benchmarking analysis

During benchmarking performance analysis several metrics have been gathered, such as:

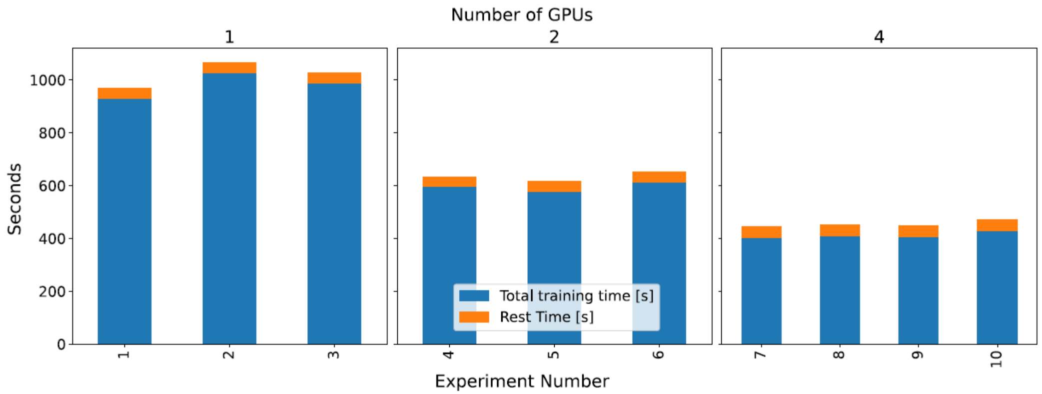

AP2 training on JSC Booster: Runtime comparison as the number of GPUs increase.

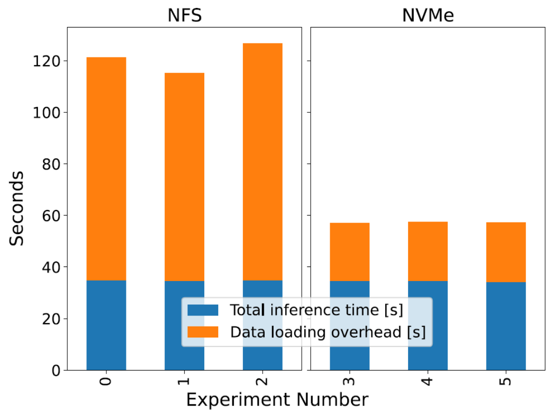

AP1 inference on E4 Intel Cluster: runtime comparison using different filesystem

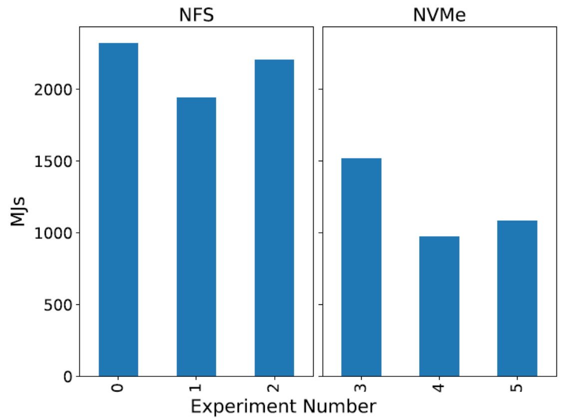

AP1 training on E4 Intel Cluster: Comparison of Action values using different filesystems

Benchmarking Metrics

Time-related

- Total runtime
- Total training time
- Training time per epoch (avg, min, max)
- Training time per iteration (avg, min, max)
- Training time of first epoch
- Model saving time

Learning-related

- Final loss (training, validation)

Energy-related

- GPU power draw (max)
- Energy consumption (GPU, node)

Image: Metrics by SBTS (up), Energy by Srinivas Agra (down) (both from NounProject.com)

Final phase of WP3

The **last benchmark** run is currently **underway**

The most **innovative hardware** has been made available to execute the applications developed in WP1:

- NVIDIA H100 GPUs
- AMD "Genoa" CPUs
- Graphcore
- Intel Sapphire Rapids CPUs
- And more…

At the end of the project, WP3 will define the **optimal compute system** design for ML applications in W&C science.

Outline for the rest of today

9:00 - 11:00 am:	Session 1 with talks to learn more about MAELSTROM
11:00 - 11:15 am:	Coffee break
11:15 - 12:45 pm:	Session 2 with talks to learn more about our EuroHPC Partner Projects
12:45 - 2:00 pm:	Lunch break
2:00 - 3:30 pm:	Session 3 with invited external speakers
3:30 - 4:00 pm:	Coffee break
4:00 - 5:30 pm:	Session 4 with invited external speakers
6:00 - ?? pm:	Drink reception

Questions?

Please get involved:
https://www.maelstrom-eurohpc.eu

If you have used MAELSTROM applications and datasets in the past, please fill in our user survey:
https://www.maelstrom-eurohpc.eu/survey

This presentation reflects the views only of the author, and the European High-Performance Computing Joint Undertaking or Commission cannot be held responsible for any use which may be made of the information contained therein.