Machine learning at ECMWF

Florian Pappenberger

Director of Forecasts & Deputy Director-General

The strength of a common goal

ECMWF Strategy: Science and technology goals for 2030

A seamless Ensemble Earth system
maximising the use of current and upcoming observations
through consistent and accurate modelling
with realistic water, energy and carbon cycles.

Use of advanced high-performance computing
big data and AI methodologies
to create a Digital Twin of the Earth
with a breakthrough in realism.

ECMWF

Earth System Science : moving forward

er enhanced NWP-Copernicus
rgies and convergence

ced model biases
ting in improved
ctions on all
cales

earth system
lexity for improved
orological
ctions
transition seasons)

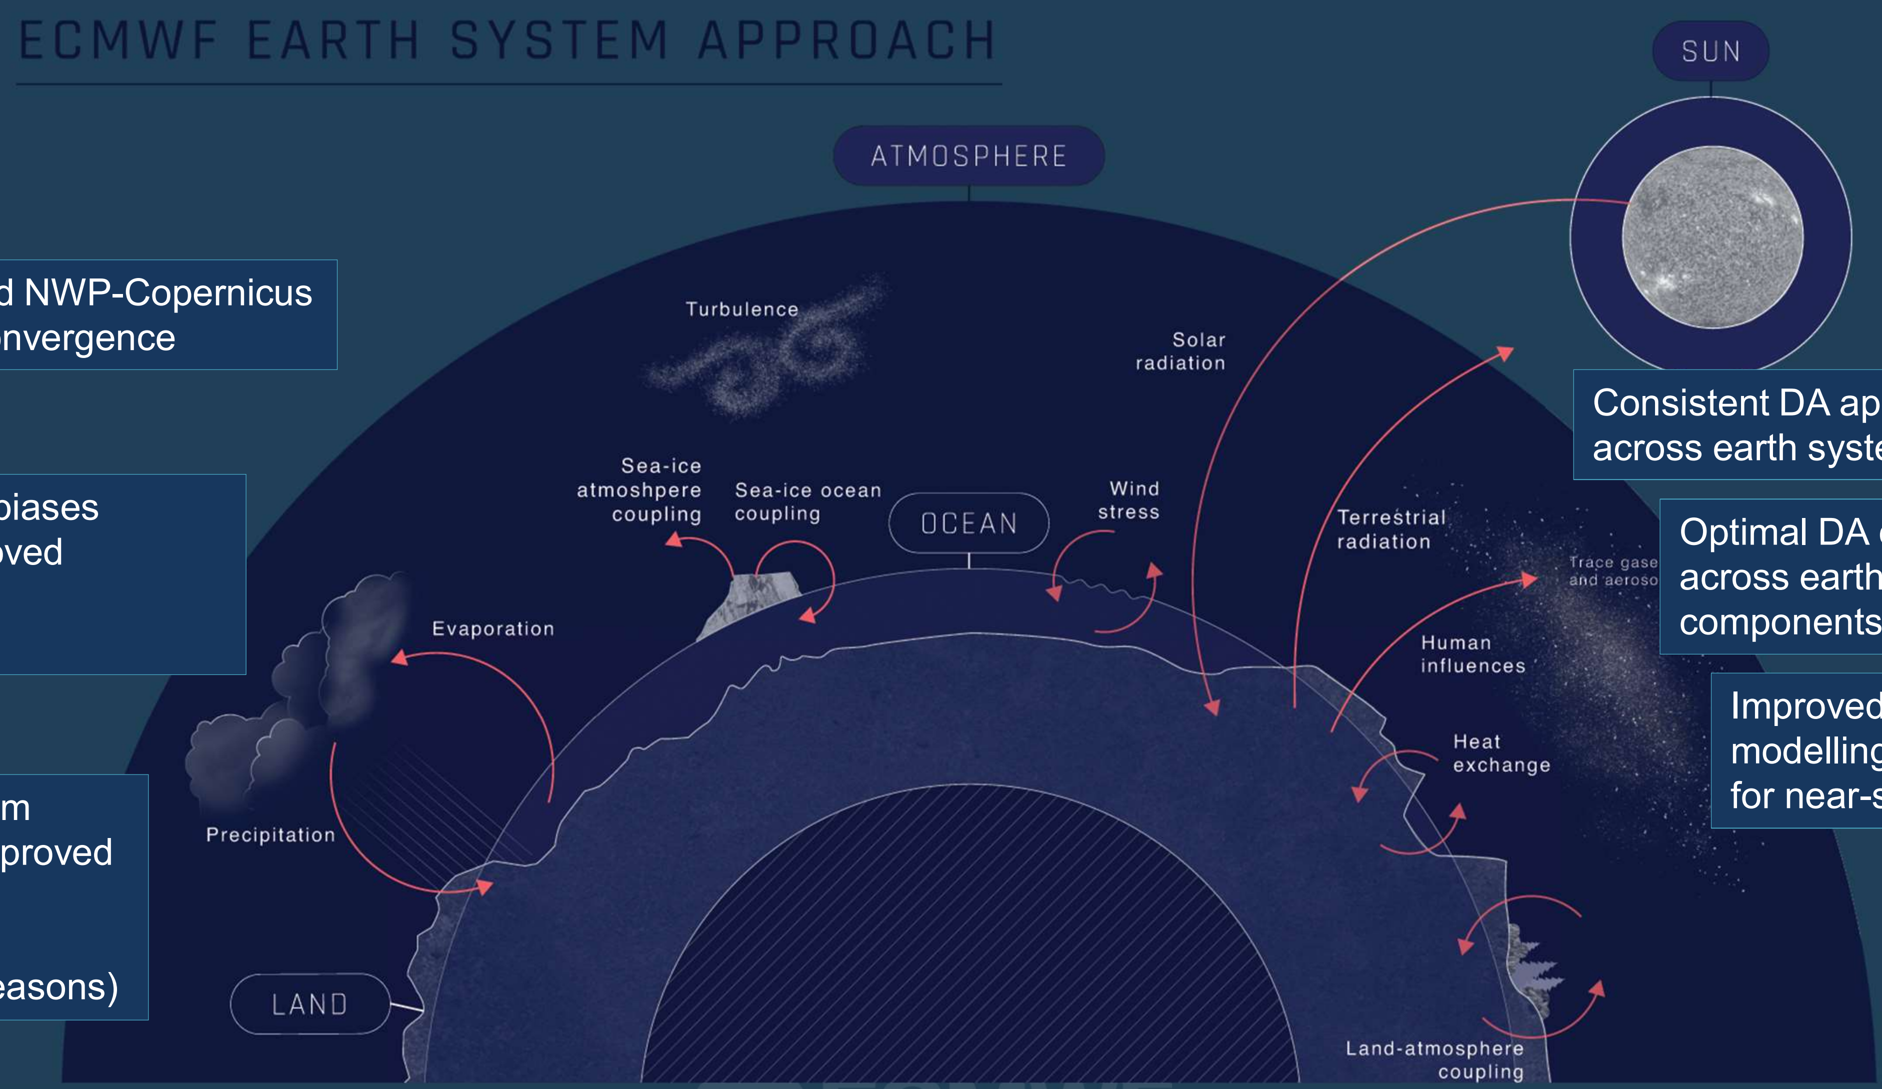

Consistent DA approaches
across earth system compon

Optimal DA coupling
across earth system
components

Improved interface
modelling
for near-surface wea

Machine Learning has been part of ECMWF forecasts for many years

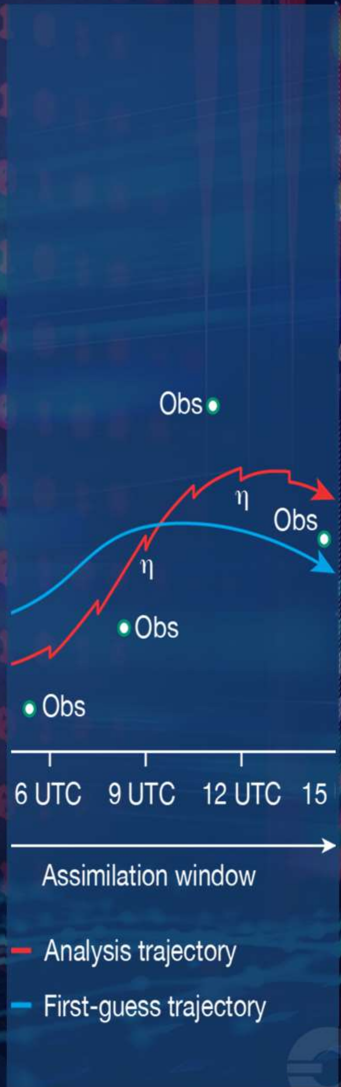

And now planning to revolutionize the full NWP workflow…

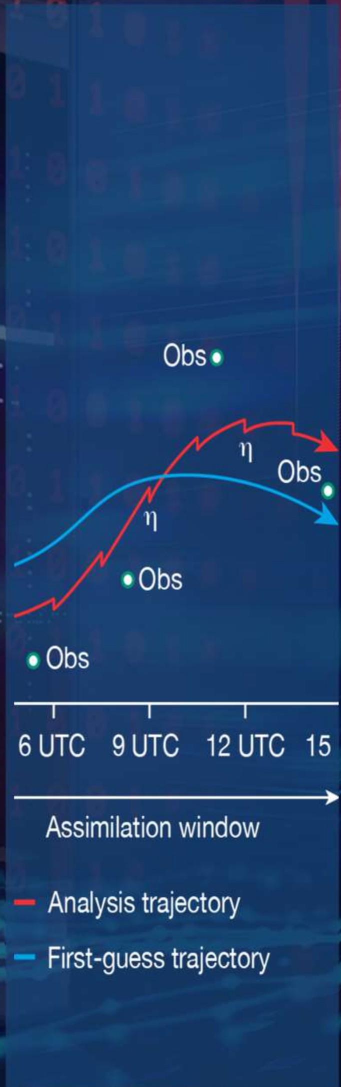

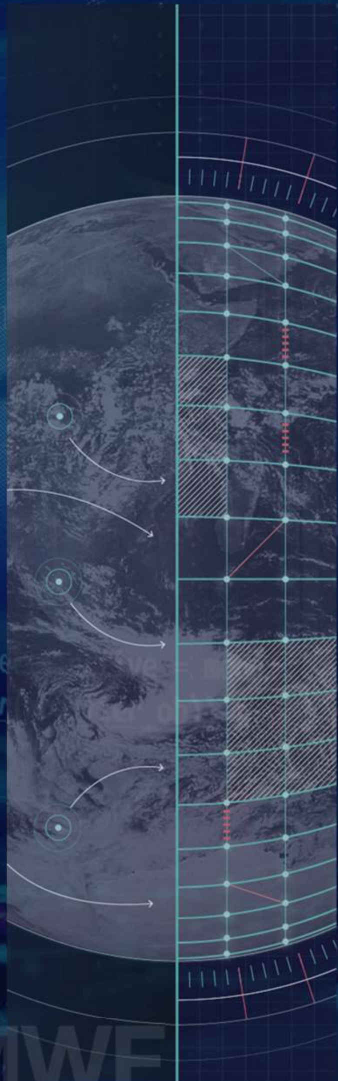

Ready for the challenge

+

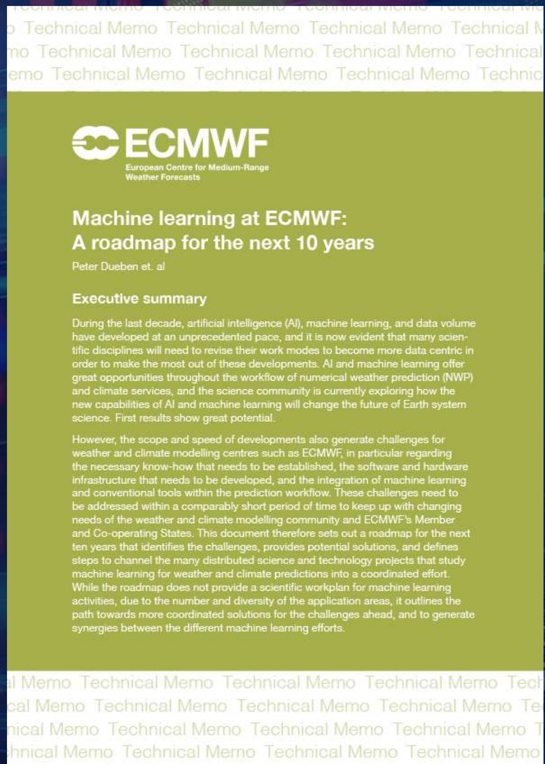
Technical Memo Technical Memo Technical Memo

ECMWF
European Centre for Medium-Range Weather Forecasts

Machine learning at ECMWF: A roadmap for the next 10 years

Peter Dueben et. al

Executive summary

During the last decade, artificial intelligence (AI), machine learning, and data volume have developed at an unprecedented pace, and it is now evident that many scientific disciplines will need to revise their work modes to become more data centric in order to make the most out of these developments. AI and machine learning offer great opportunities throughout the workflow of numerical weather prediction (NWP) and climate services, and the science community is currently exploring how the new capabilities of AI and machine learning will change the future of Earth system science. First results show great potential.

However, the scope and speed of developments also generate challenges for weather and climate modelling centres such as ECMWF, in particular regarding the necessary know-how that needs to be established, the software and hardware infrastructure that needs to be developed, and the integration of machine learning and conventional tools within the prediction workflow. These challenges need to be addressed within a comparably short period of time to keep up with changing needs of the weather and climate modelling community and ECMWF's Member and Co-operating States. This document therefore sets out a roadmap for the next ten years that identifies the challenges, provides potential solutions, and defines steps to channel the many distributed science and technology projects that study machine learning for weather and climate predictions into a coordinated effort. While the roadmap does not provide a scientific workplan for machine learning activities, due to the number and diversity of the application areas, it outlines the path towards more coordinated solutions for the challenges ahead, and to generate synergies between the different machine learning efforts.

Technical Memo Technical Memo Technical Memo

→

CMWF’s machine learning roadmap

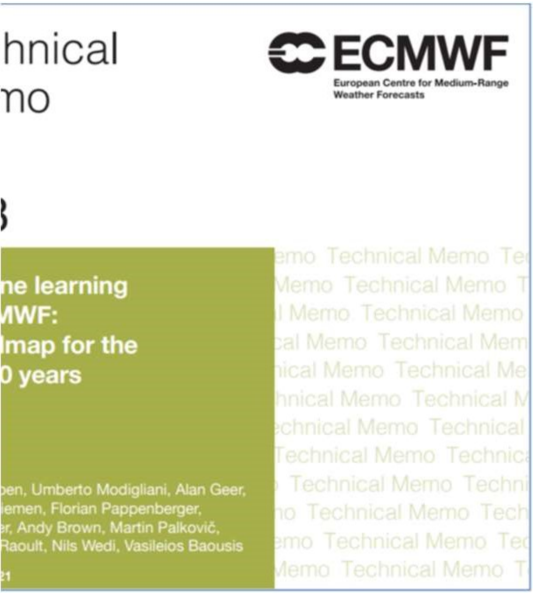

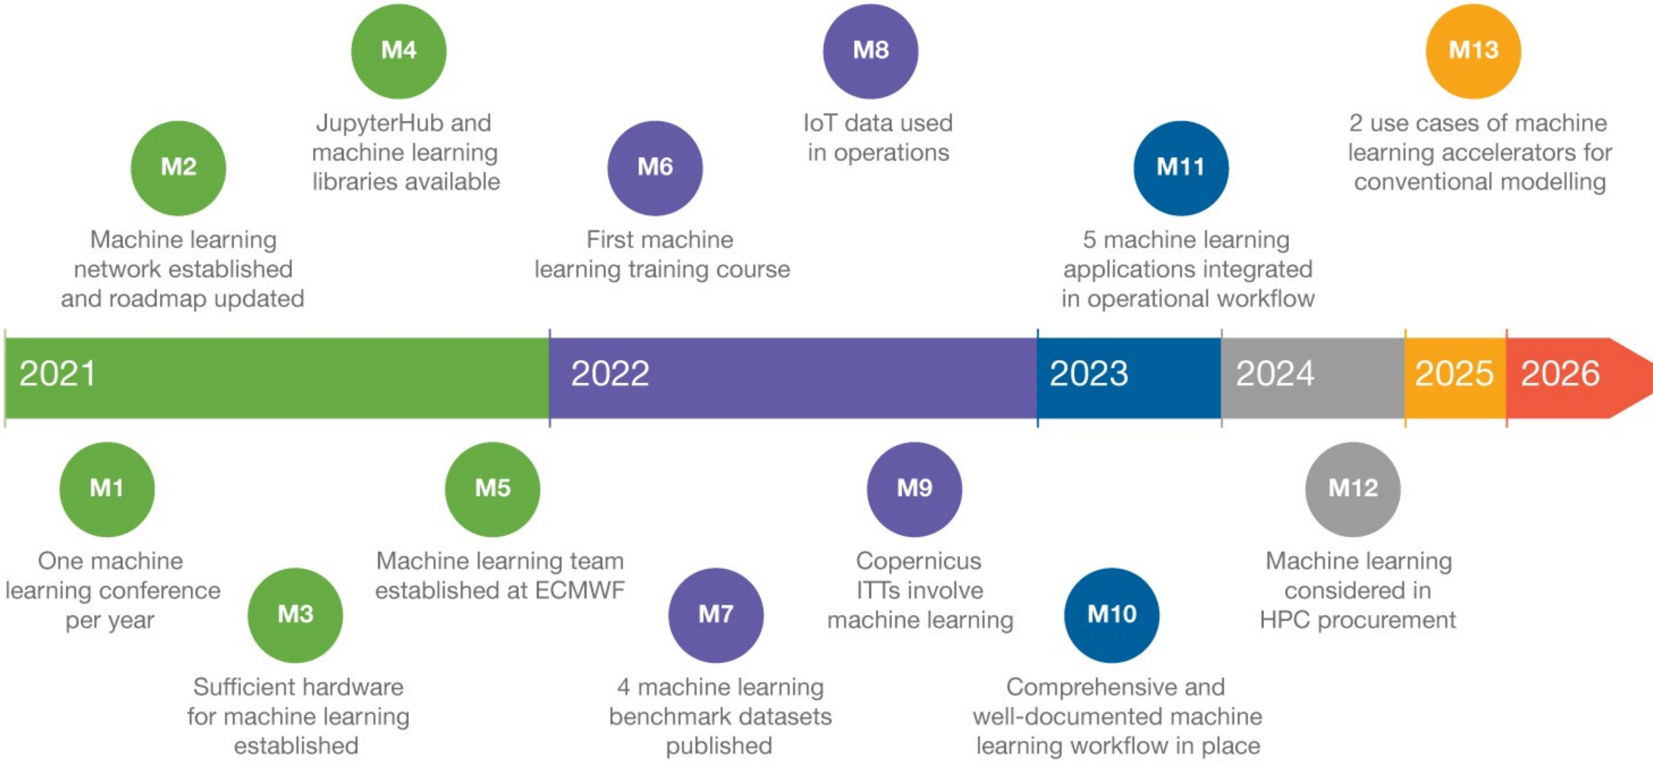

upporting Roadmap I
entre of Excellence (COE) in Weather & Climate Modelling

ECMWF

Atos NVIDIA

Aim

- develop new techniques to support next-generation weather forecasting
- help boost climate and weather discovery and innovation
- prepare ECMWF for future HPC and data handling architectures.

pporting Roadmap II: Infrastructure

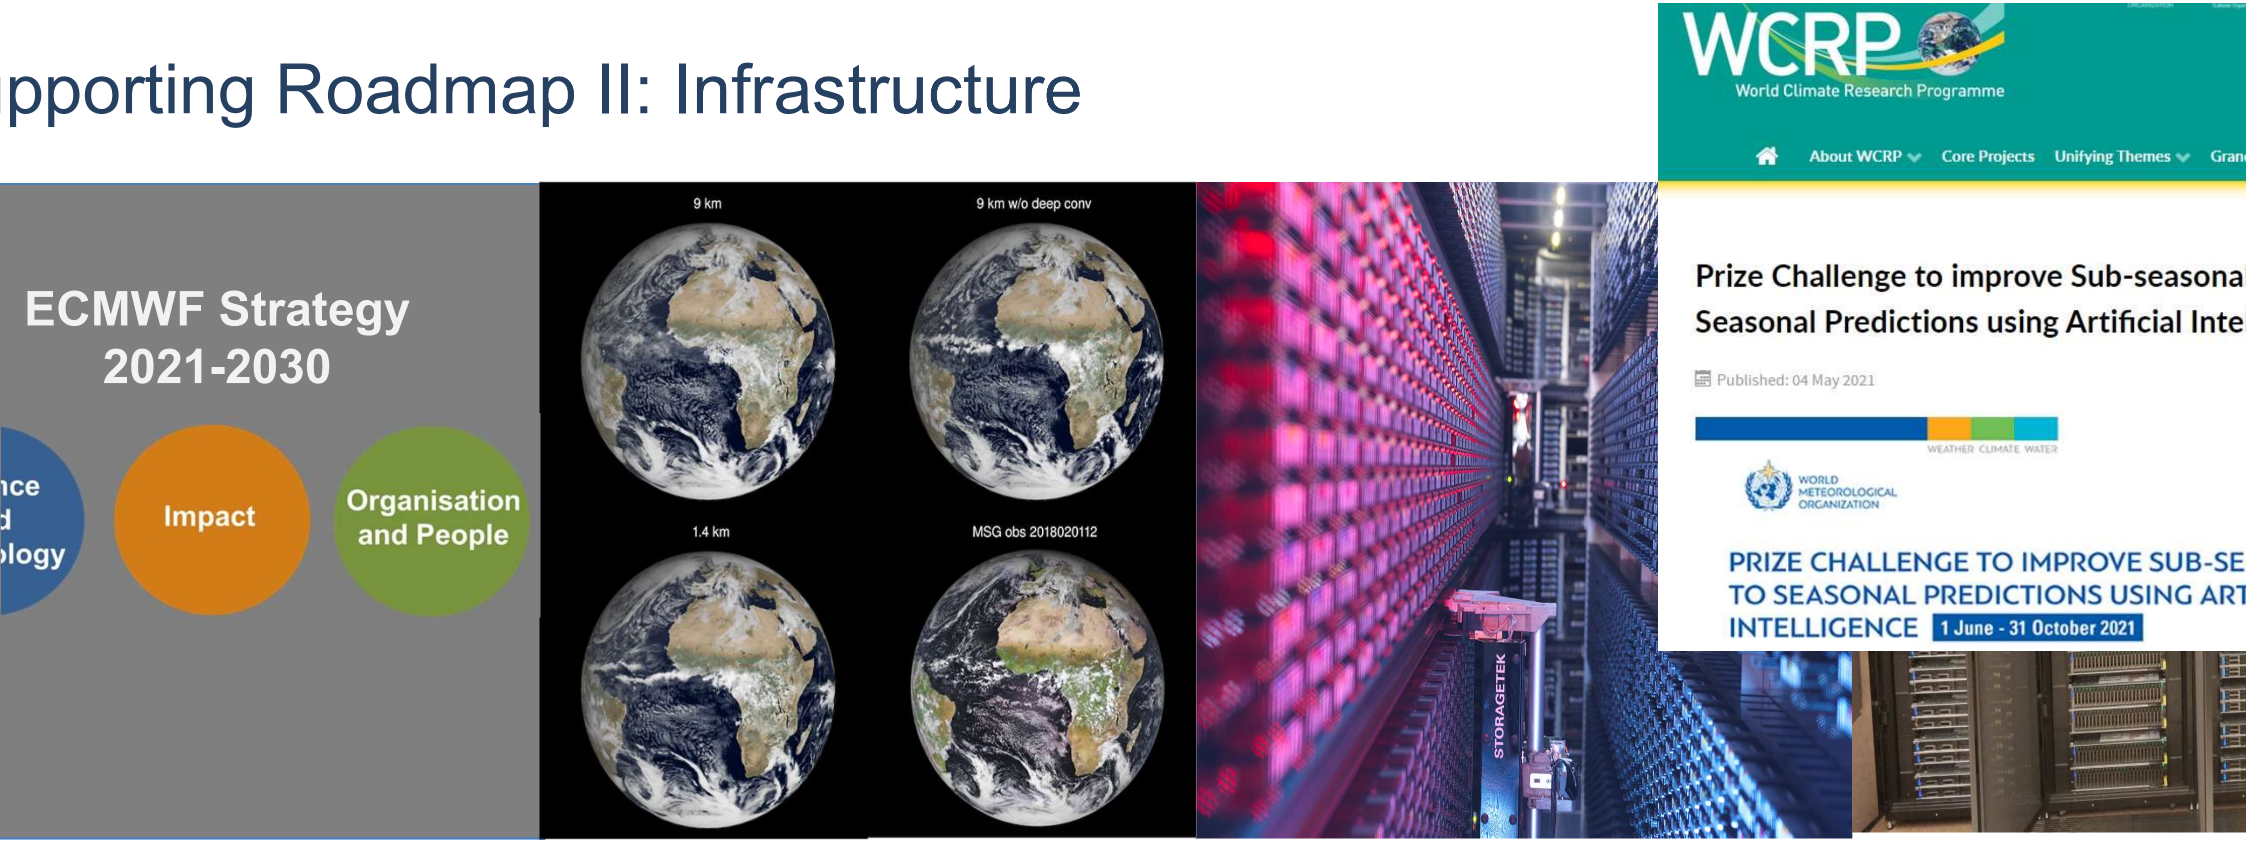

The European Weather Cloud aims to become the ***cloud-based collaboration platform for meteorological application development & operations in Europe*** *and contributes to the digital transformation of the European Meteorological Infrastructure*

"a **community** cloud"

ECMWF EUMETSAT EUROPEAN WEATHER CLOUD
CLOUD COMPUTING-BASED INFRASTRUCTURE FOCUSED ON THE NEEDS OF THE METEOROLOGICAL COMMUNITY

upporting Roadmap IIIa: H2020
e MAELSTROM project

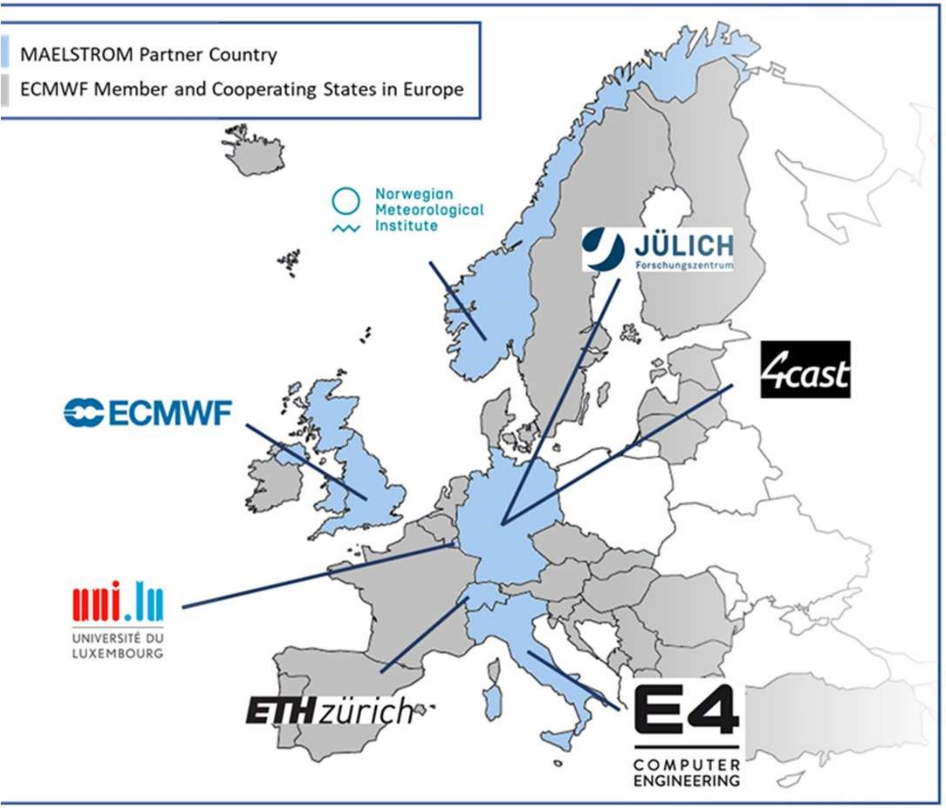

Horizon 2020 Programme

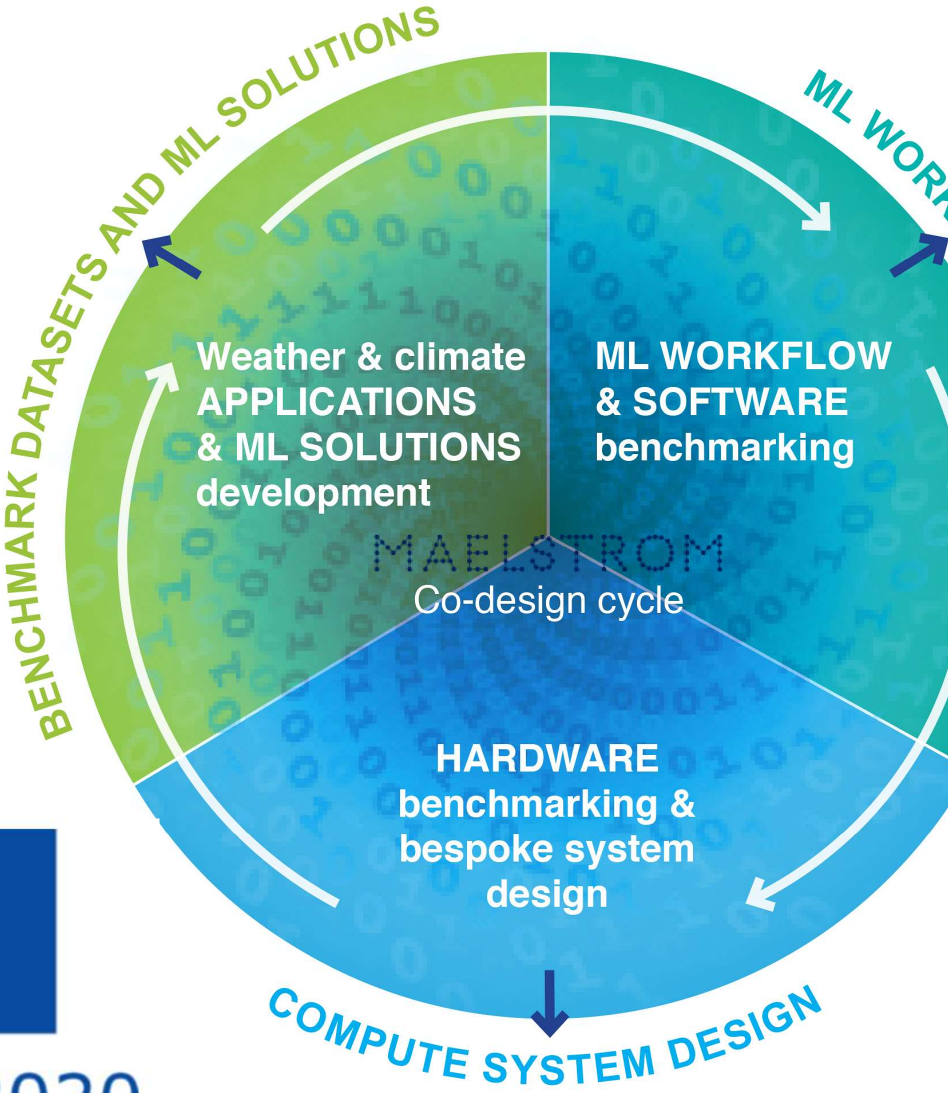

e first **benchmark datasets** have been published!

pporting Roadmap IIIb: H2020 & Partnerships
wards pre-operational machine learning tropical cyclone detection

- ECMWF, NOAA and NVIDIA collaboration

Cyclone tracks
from *"IB-tracks"*

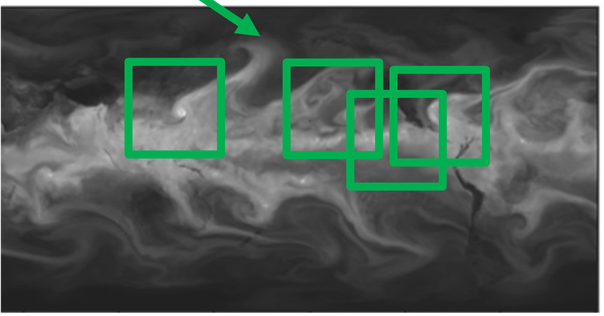

orecast (or satellite) fields

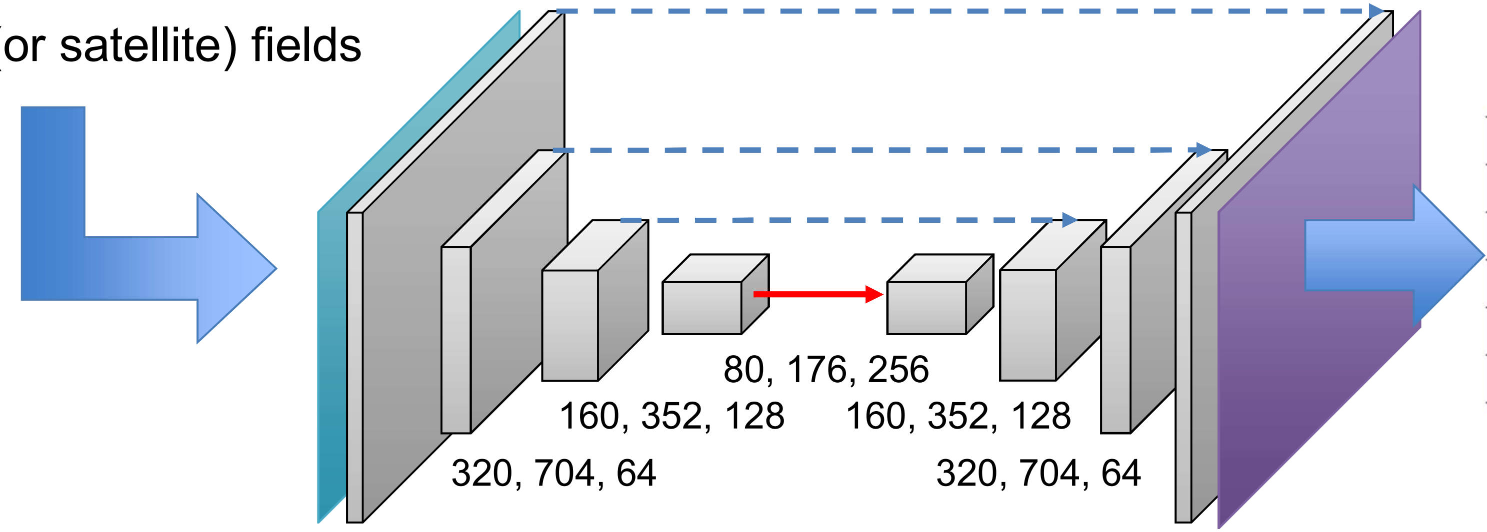

Horizon 2020
Programme

Andrea Castelletti
@hydroaholics

Proudly announcing that our H2020 project p
CLINT (CLImate INTelligence) has been funde
will work with a unique group of scientists on i
climate science and services using #Machine
in a number of climate hotspots

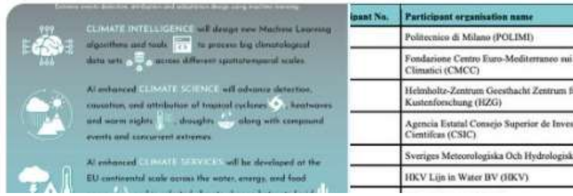

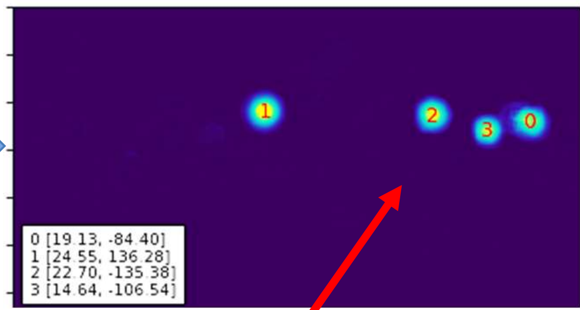

Predicted
Cyclone centers

ECMWF

upporting Roadmap IV: Innovation Programme
uropean Summer of Weather Code

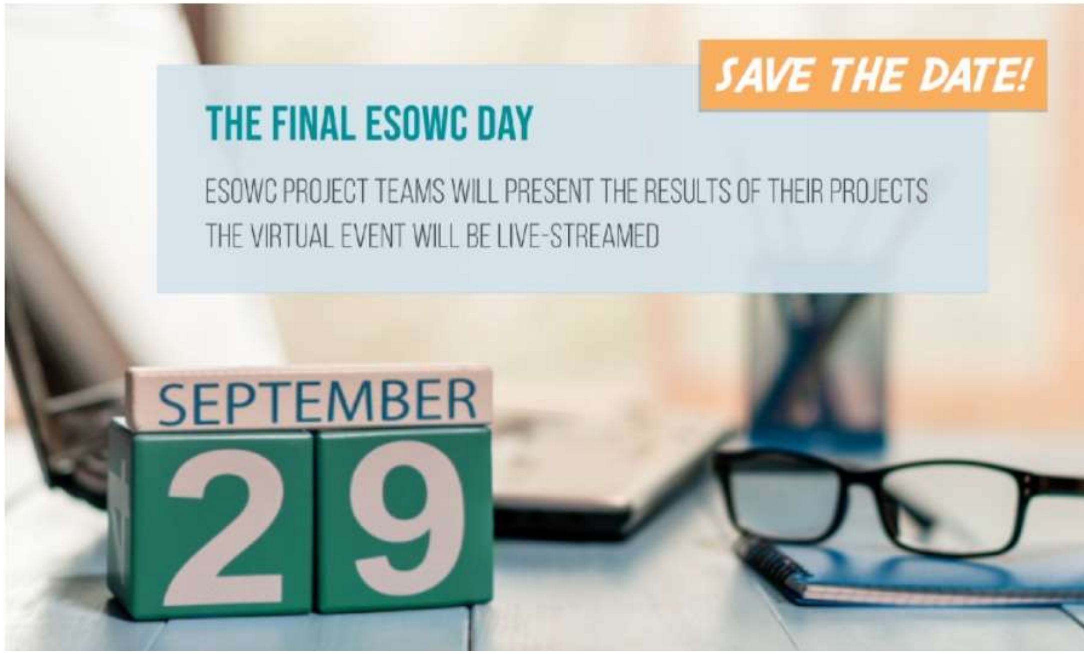

MaLePoM
(Machine Learning for Pollution Monitoring)

CliMetLab - Machine Learning on weather and climate dat

fridge.jl: Compressing atmospheric data into its real information content

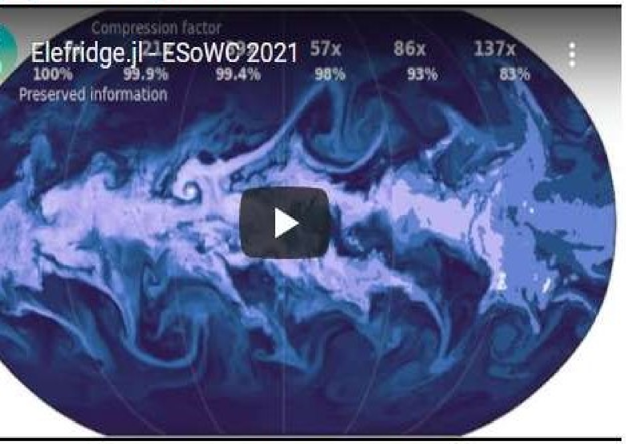

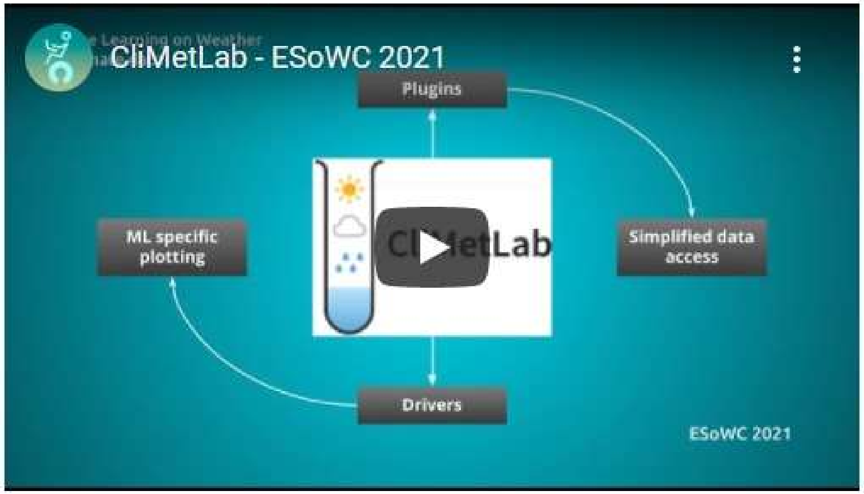

upporting Roadmap V: YOU

Machine learning for numerical weather predictions and climate services – A workshop for Member and Co-operating States

| 14-16 April 2021

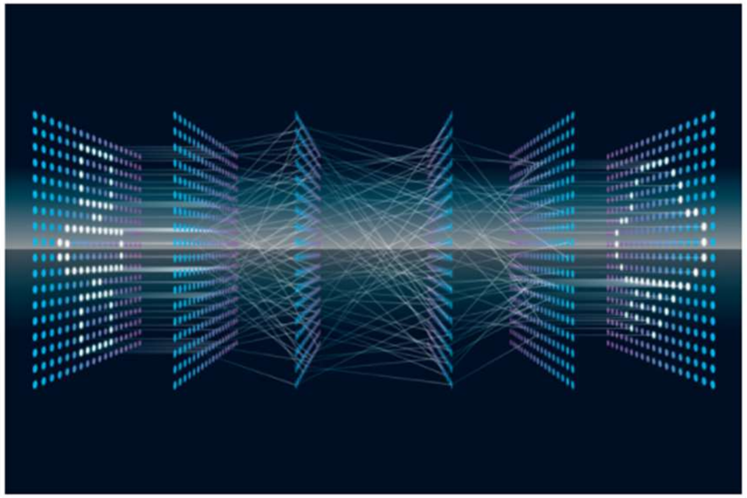

Workshop overview

This virtual workshop aimed to update ECMWF's Member and Co-operating States about current machine learning efforts at ECMWF and to allow for the active involvement in the realisation of ECMWF's machine learning roadmap. The workshop allowed for active

Application for Member State short-term secondment to ECMWF

ECMWF invites short-term secondments from Member State and Co-operating State hydro-meteorological institutes.

Projects can cover all areas of work, typically science, forecast delivery, computing, environmental applications, administration and communication. Any secondment proposal must be agreed with your line management.

ECMWF can offer partial funding to support such stays. The work arrangements can be any period from several weeks up to a maximum of three months, either as a continuous stay or a sequence of shorter stays.

The seniority of the candidate can be at any level, from trainee to experienced staff.

Member State workshops organized by Peter Dueben (all questions to him!)

Check website under Jobs

ECMWF

atus of machine learning at ECMWF

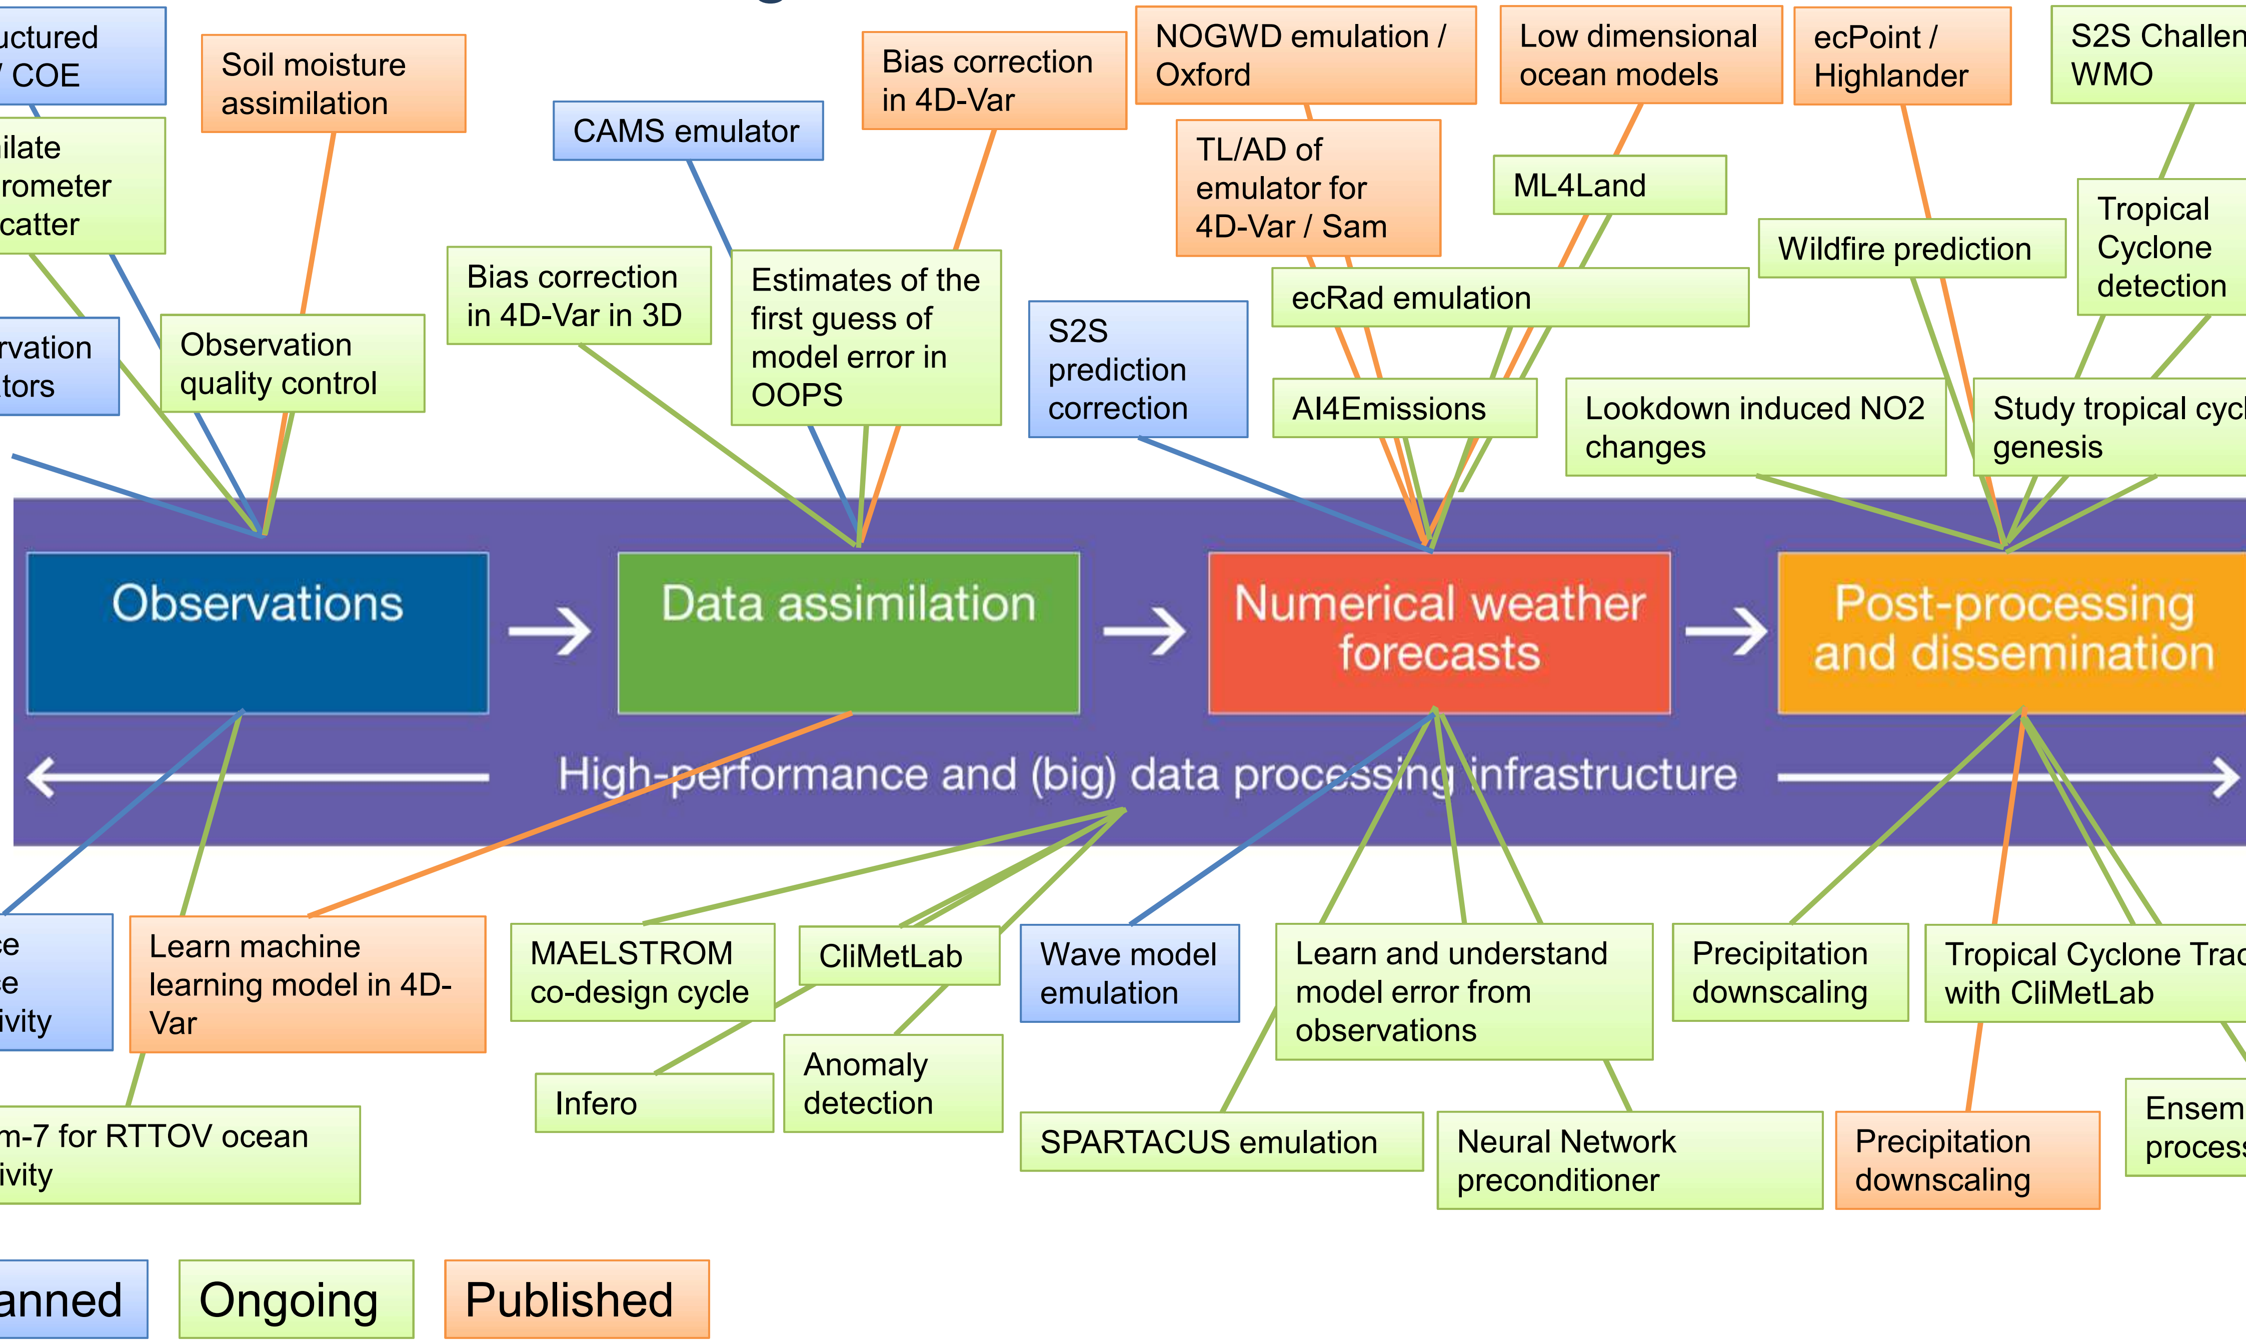

emulate the 3D cloud effect in radiation

epresent 3D cloud effects for radiation (SPARTACUS) within simulations of the Integrated Forecast Model
time slower than the standard radiation scheme (Tripleclouds)

we emulate the difference between Tripleclouds and SPARTACUS using neural networks?

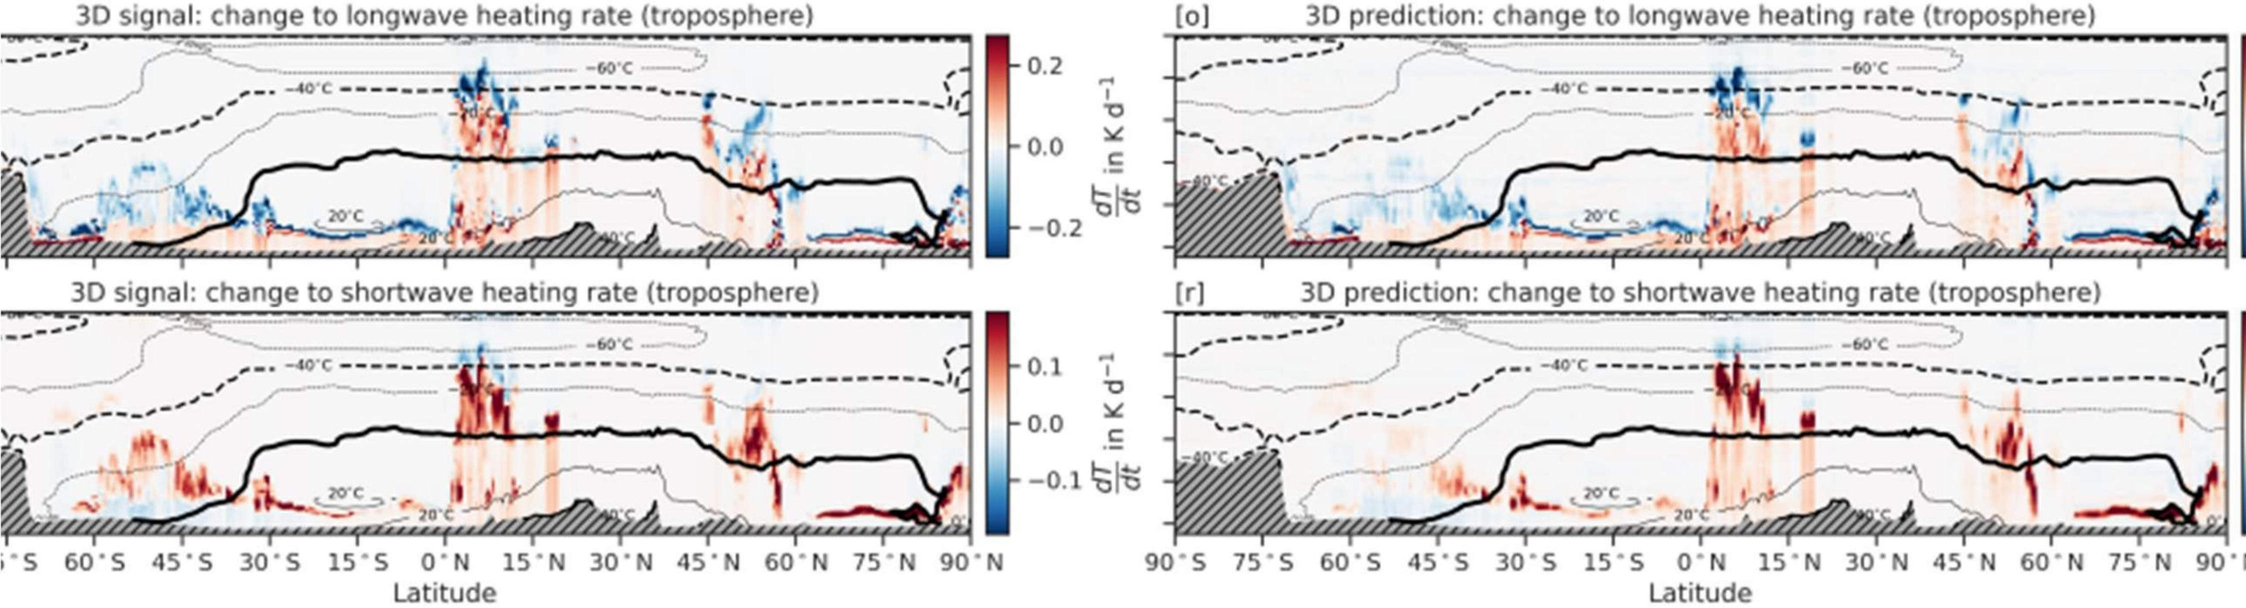

	Tripleclouds	SPARTACUS	Neural Network	Tripleclouds+Neural Network
ative Cost	1.0	4.4	0.003	1.003

Machine learning applied to forecast 2m temperature and 10m wind

Fenwick Cooper, Zied Ben Bouallegue, Matthew Chantry, Peter Düben, Peter Bechtold, Irina Sandu

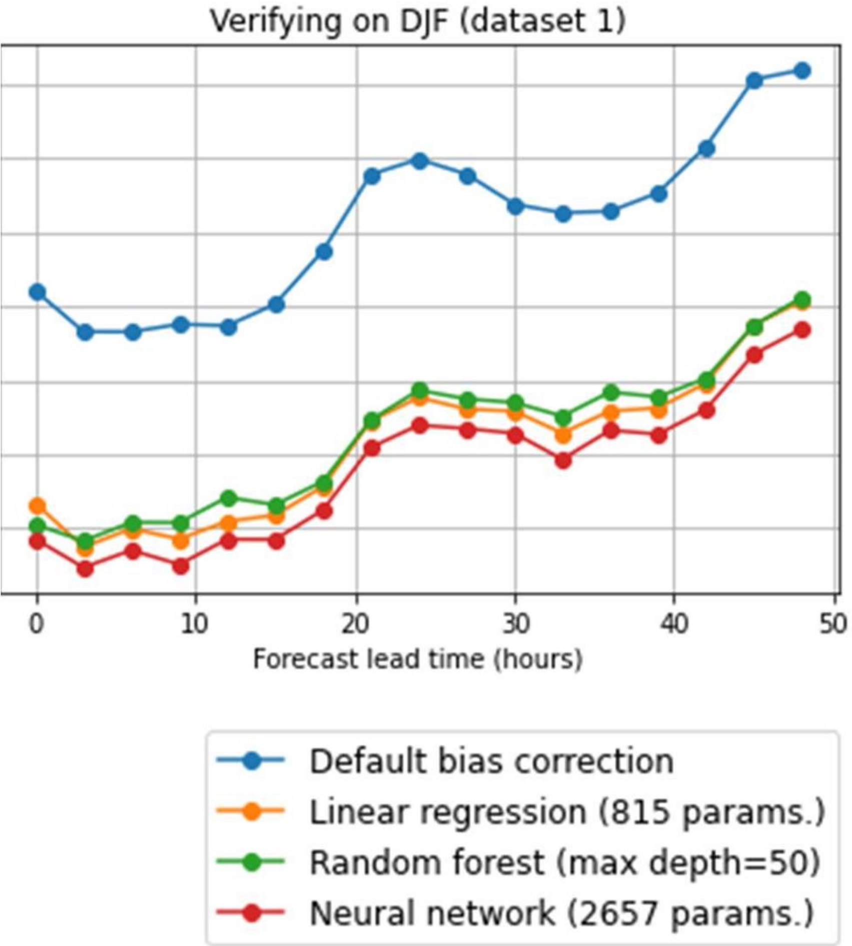

Example: 2m temperature, Winter 2020, 1 year of training data

Root-Mean-Squared Error (RMSE) with respect to station measuremen
All stations (left) – Individual stations change (below)

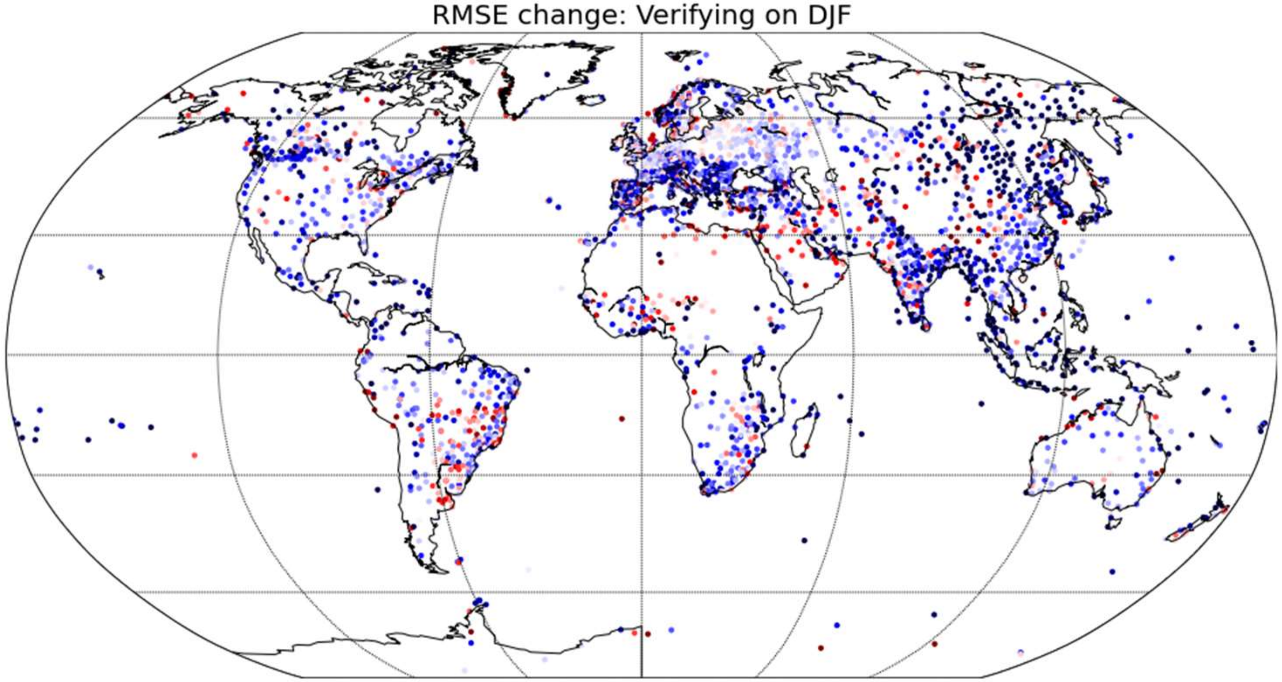

W

°C

Be

ECMWF | EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS

The strength of a common goal

Machine learning for weather predictions

Peter Dueben

Royal Society University Research Fellow & ECMWF's Coordinator for Machine Learning and AI Activities

The strength of a common goal

arch used resources of the Oak Ridge
Computing Facility (OLCF), which is a DOE
cience User Facility supported under Contract
00OR22725.

The ESIWACE, MAELSTROM and AI4Copernicu
have received funding from the European Un
grant agreement No 823988, 955513 and 10101

Let's start with definitions

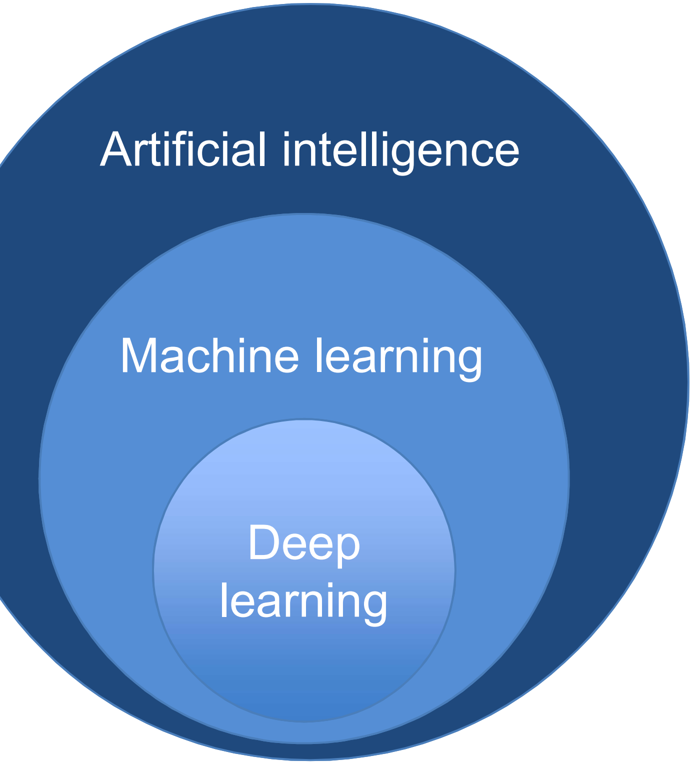

Artificial intelligence (AI) *is intelligence demonstrated by machines, i*
contrast to the natural intelligence displayed by humans (Wikipedia)
Example: A self-driving car stops as it detects a cyclist crossing

Machine learning (ML) *is the scientific study of algorithms and statisti*
models that computer systems use to perform a specific task without u
explicit instructions… (Wikipedia)
Example: To learn to distinguish between a cyclist and other things fro

Deep learning *is part of a broader family of machine learning methods*
based on artificial neural networks (Wikipedia)
Example: The technique that is used to detect a cyclist in a picture

Deep learning and artificial neural networks as one example of machine learning

The concept:
Take input and output samples from a large data set
Learn to predict outputs from inputs
Predict the output for unseen inputs

The key:
Neural networks can learn a complex task as a “black box”
No previous knowledge about the system is required
More data will allow for better networks

The number of applications is increasing by day:
Image recognition
Speech recognition
Healthcare
Gaming
Finance
Music composition and art
…

And weather?

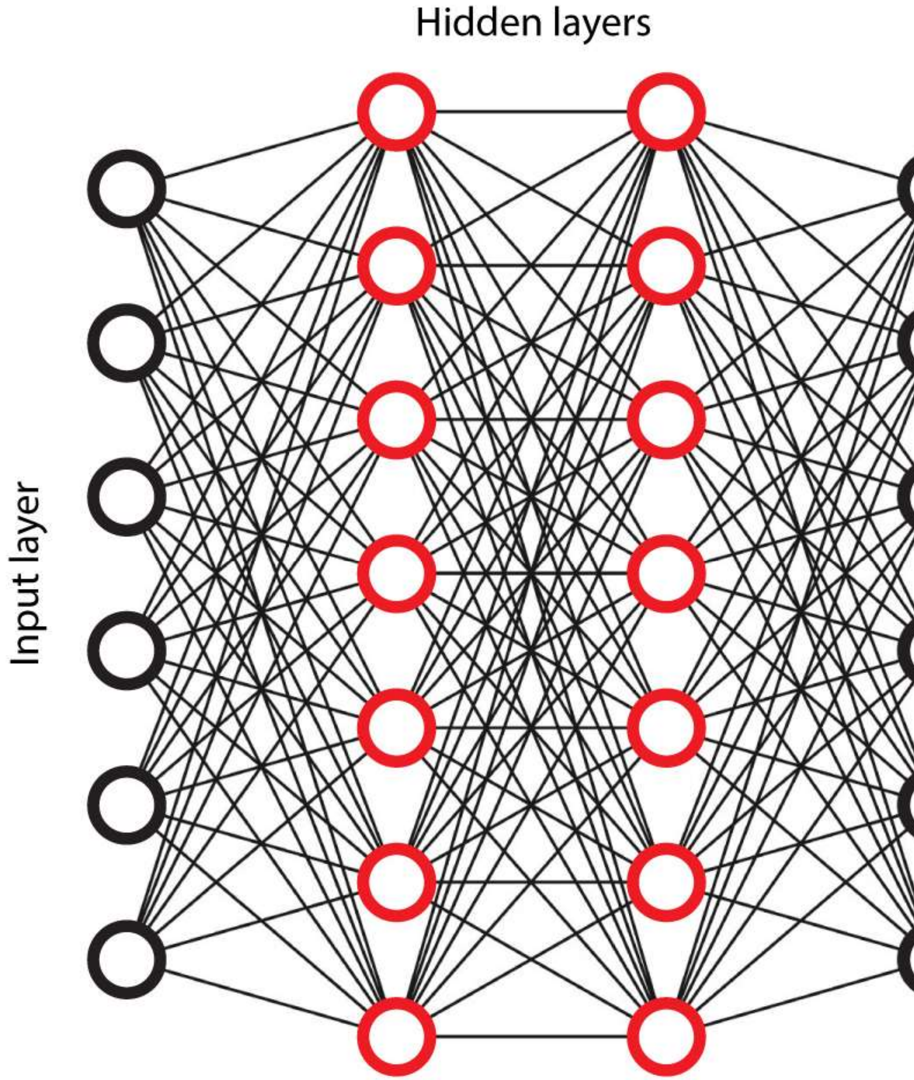

Decision trees and random forrests

Decisions fork in tree structures until a prediction is made.

"Random forest" methods are training a multitude of decision trees using a mean predictions or the value with the most hits as a result.

Decision trees are often fast and accurate and they are able to conserve some of the properties of the system.

Decision trees often require a lot of memory (as they serve as an efficient look-up table).

An example for ecPoint:

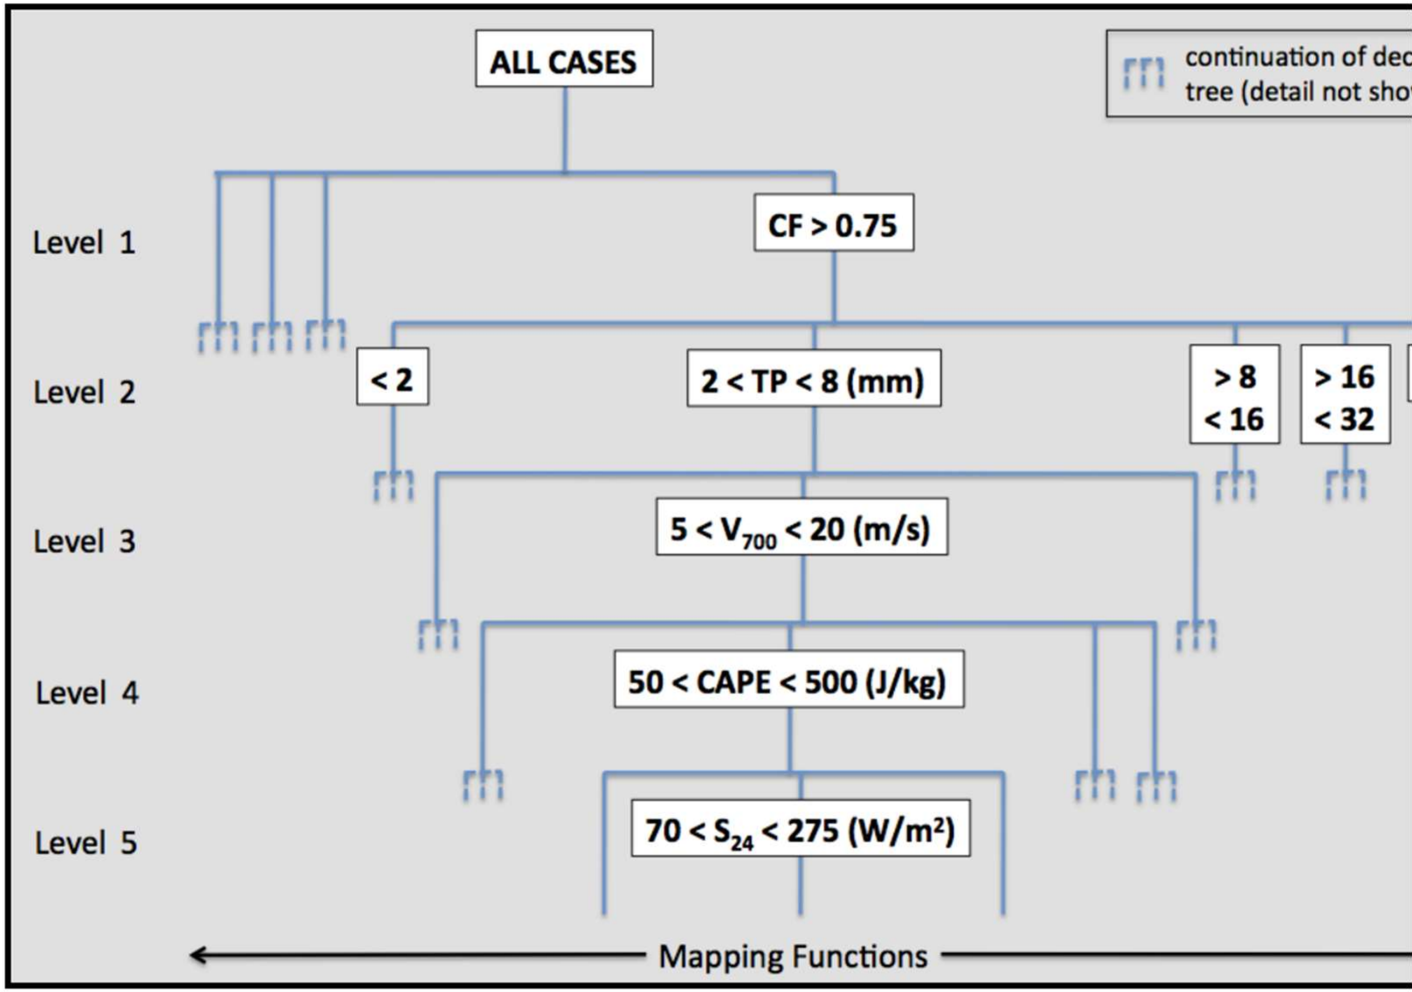

Hewson and Pillosu 2020

Two families of machine learning

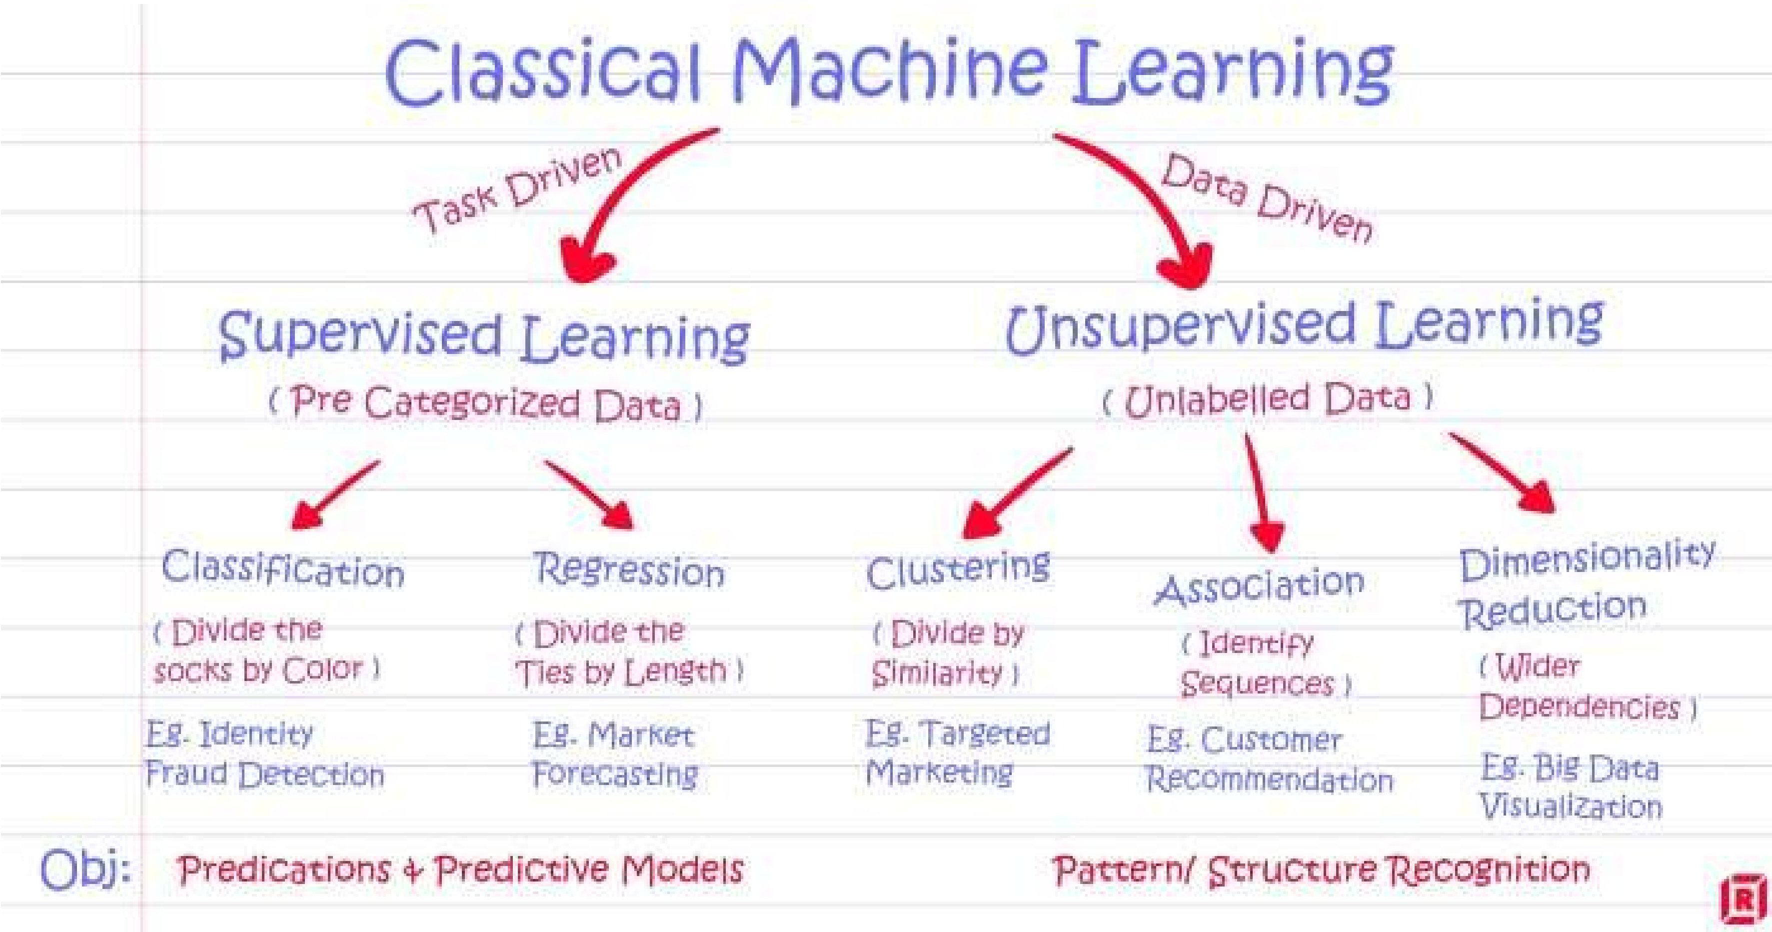

Source: https://medium.com/@recrosoft.io/supervised-vs-unsupervised-learning-key-differences-cdd46206cdcb

Why would machine learning help in weather and climate predictions?

ictions of weather and climate are difficult:

he Earth is huge, resolution is limited and we cannot represent l important processes within model simulations

he Earth System shows “chaotic” dynamics which makes it fficult to predict the future based on equations

l Earth System components (atmosphere, ocean, land surface, oud physics,…) are connected in a non-trivial way

ome of the processes involved are not well understood

ever, we have a huge number of observations and Earth em data

here are many application areas for machine learning in umerical weather predictions

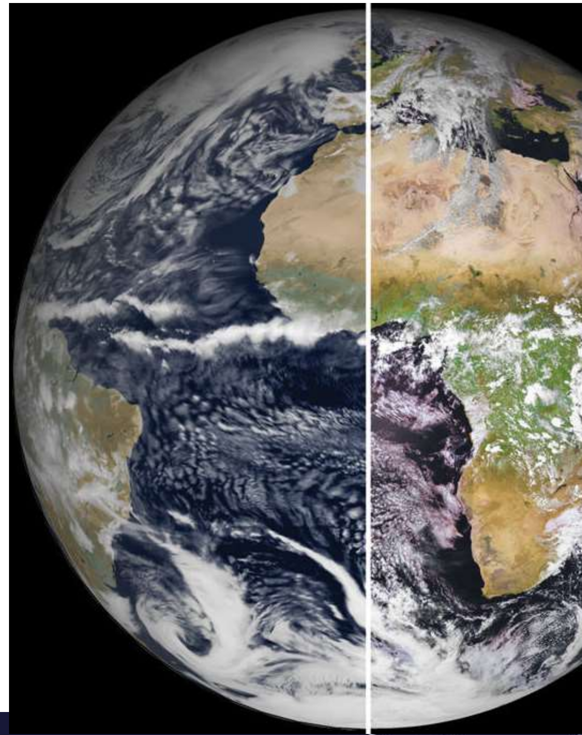

...y is machine learning so hip at the moment?

- ...ncrease in data volume
- ...New computing hardware
- ...New machine learning software
- ...ncrease in knowledge

Slide from Torsten Hoefler (ETH)

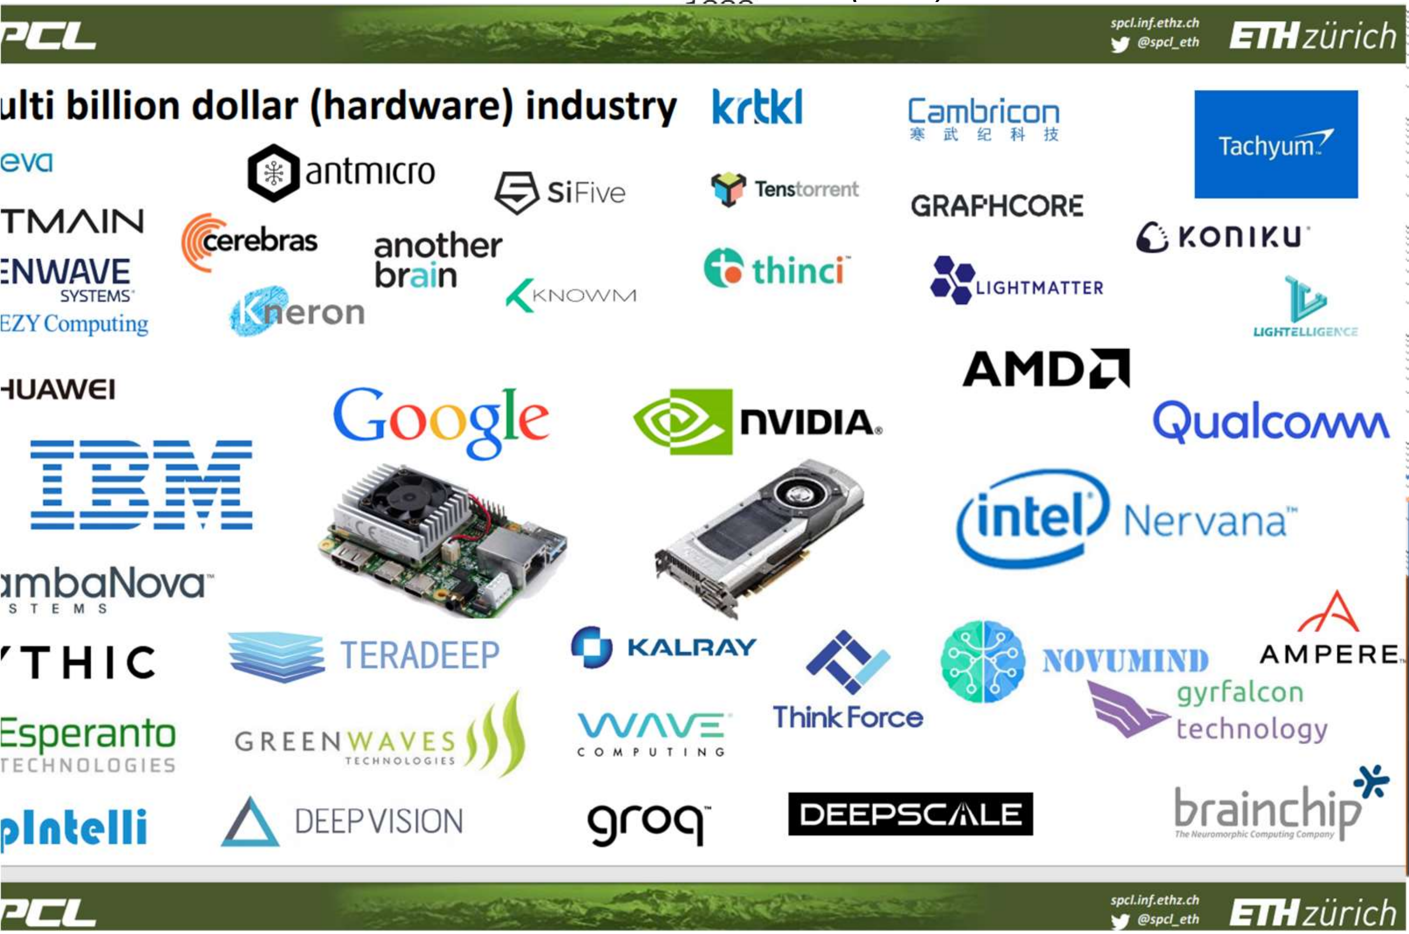

Bauer et al. ECMWF SAC paper 2019

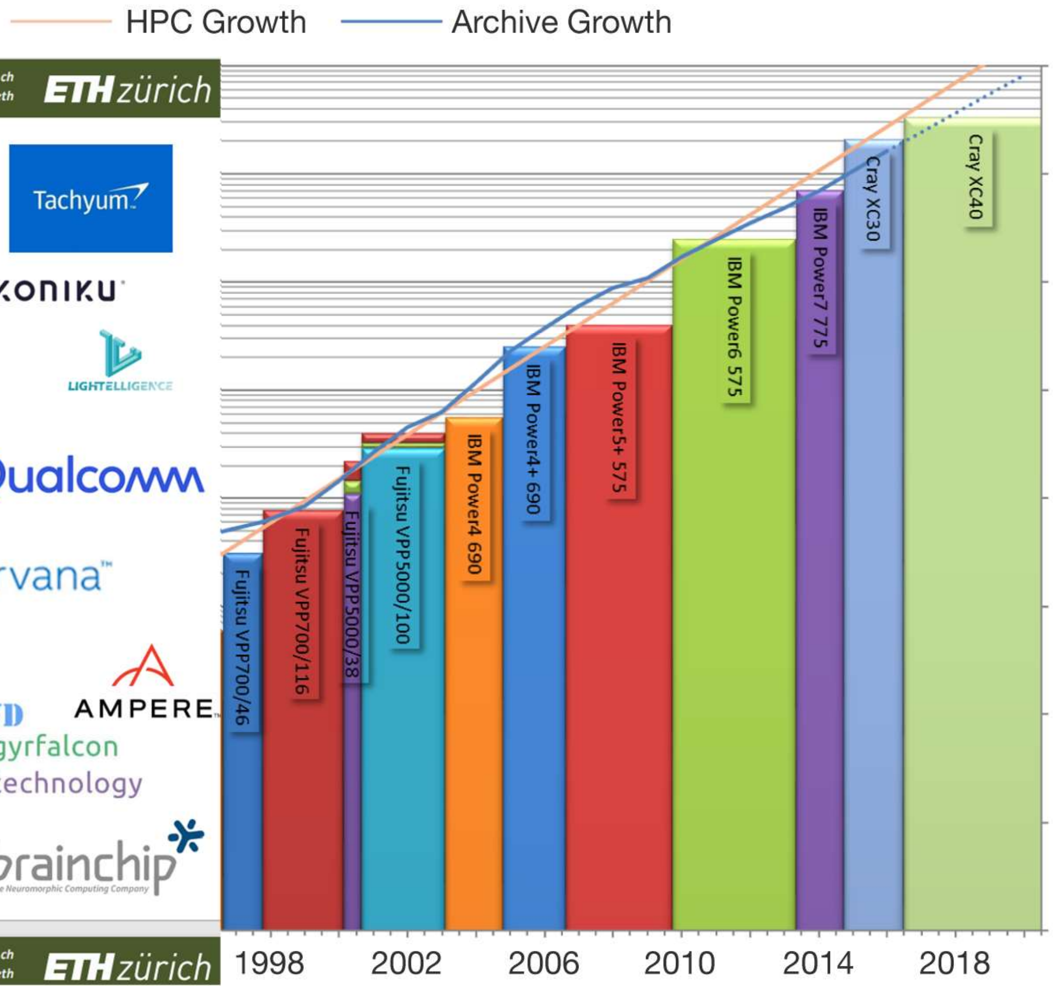

hat will machine learning for numerical weather and climate predicti look like in 10 years from now?

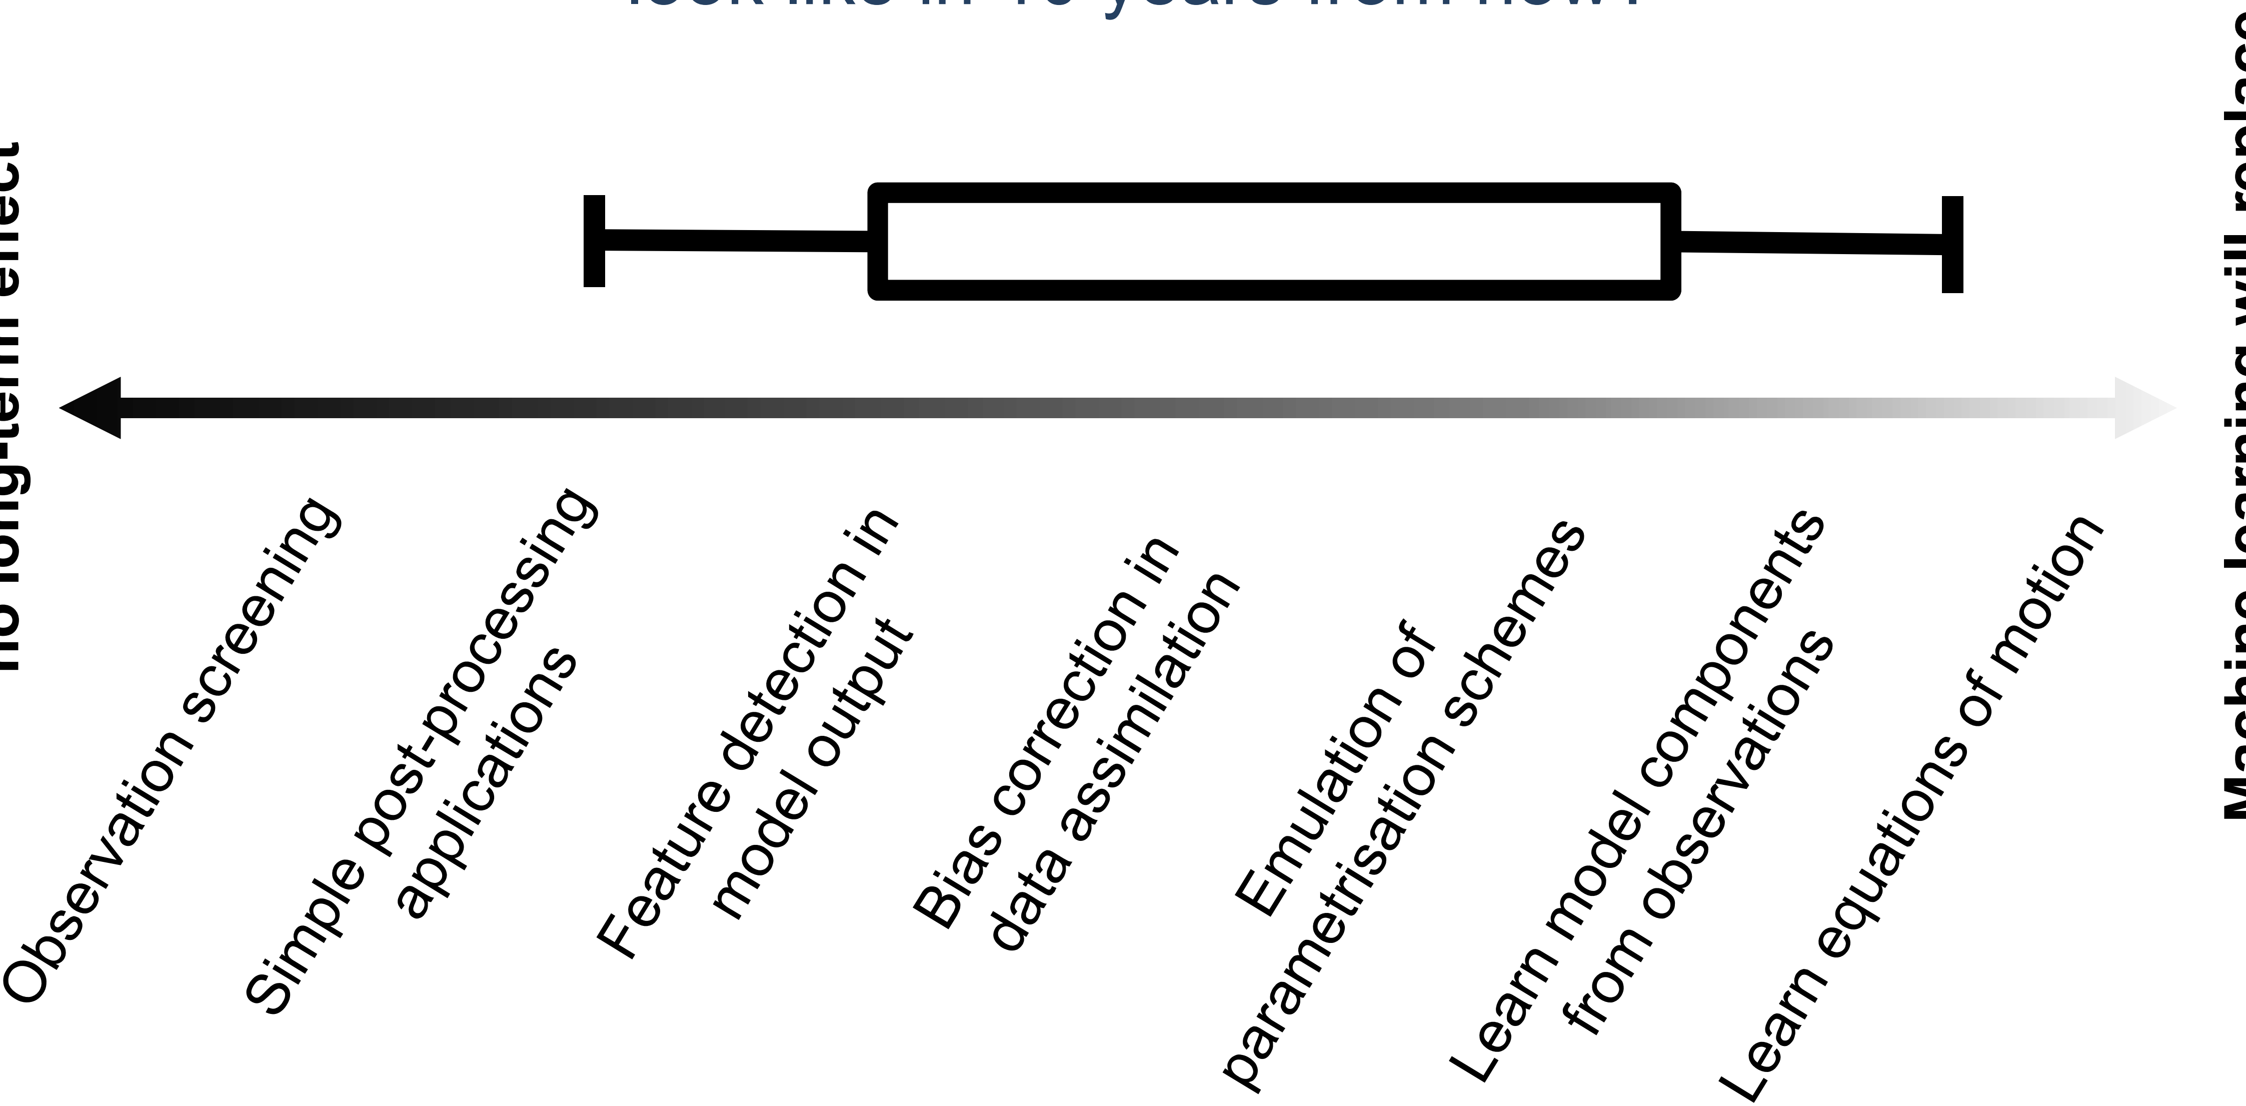

e uncertainty range is still very large...

an we replace conventional weather forecast systems by deep learni

e could base the entire model on neural networks and trash the conventional models.?
ere are limitations for existing models and ECMWF provides access to hundreds of petabytes of

simple test configuration:

- We retrieve historical data (ERA5) for geopotential at 500 hPa (Z500) for the last decades (>65,000 global data sets)
- We map the global data to a coarse two-dimensional grid (60x31)
- We learn to predict the update of the field from one hour to the next using deep learning
- Once we have learned the update, we can perform predictions into the future

physical understanding is required!

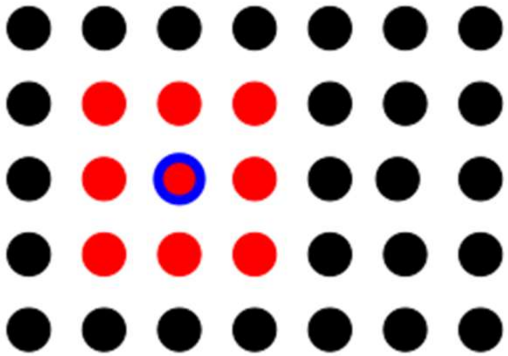

Dueben and Bauer GMI

an we replace conventional weather forecast systems by deep learnin

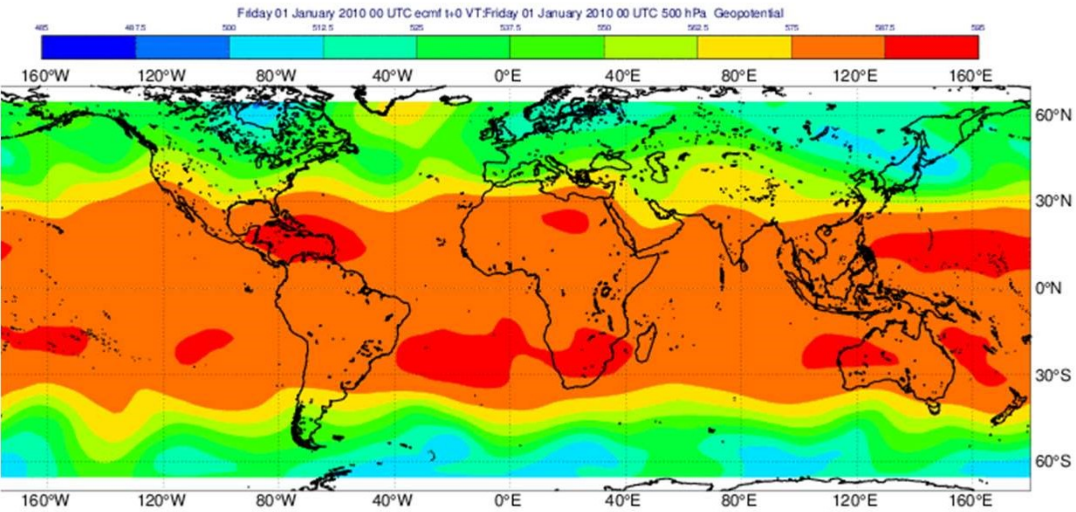

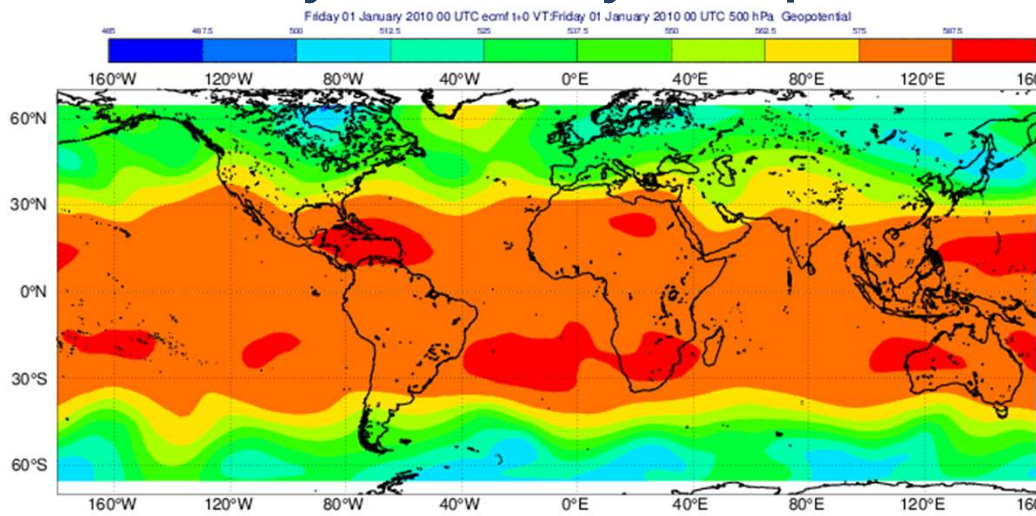

e evolution of Z500 for historic data and a neural network prediction.
you tell which one is the neural network?

- The neural network is picking up the dynamics nicely.
- Forecast errors are comparable if we compare like with like.
- There is a lot of progress at the moment.
  cher and Messori GMD 2019; Weyn, Durran, and Caruana JAMES 2019; Rasp and Thuerey 2020...
- s this the future for medium-range weather predictions?

kely…

- The simulations change dynamics in long integrations and it is unclear how o fix conservation properties.
- t is unknown how to increase complexity and how to fix feature interactions.
- There are only ~40 years of data available.

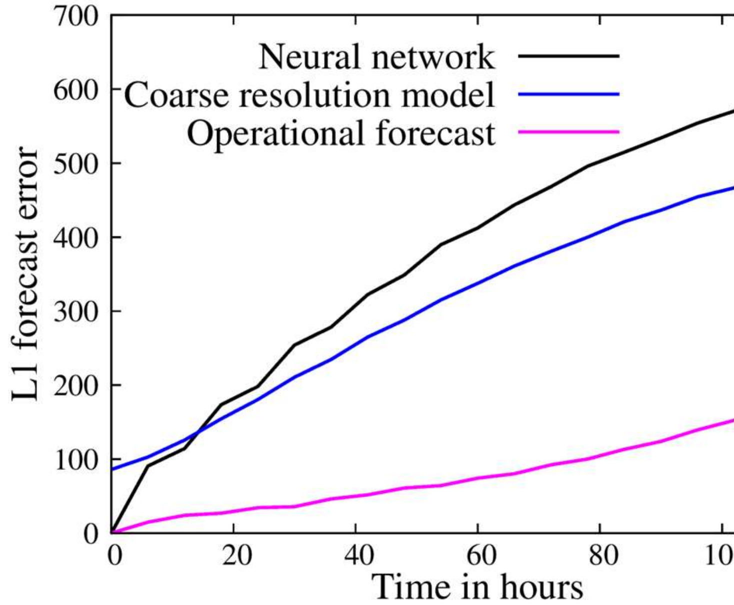

Dueben and Bauer GMI

Can we replace conventional Earth System models by deep learning

our MetNET precipitation predictions by Google:
awal, Barrington, Bromberg, Burge, Gazen, Hickey arXiv:1912.12132

NOAA forecast Ground truth Machine learning:

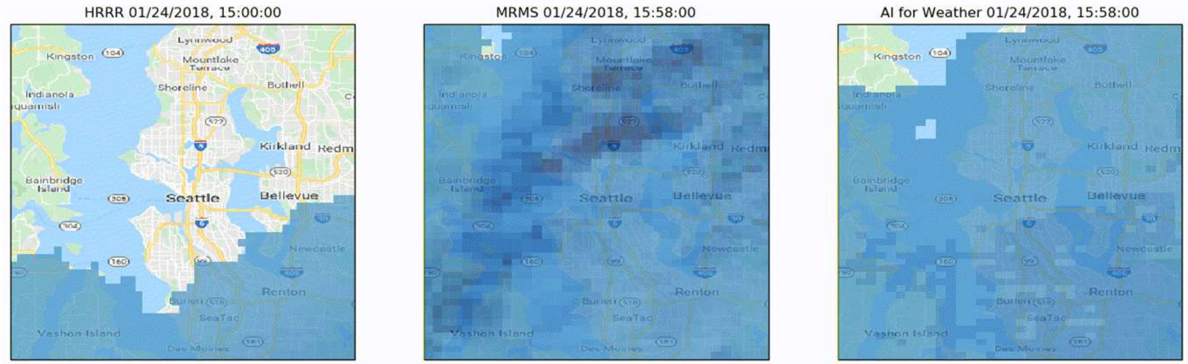

p learning for multi-year ENSO forecasts: Ham, Kim, Luo *Nature* 2019

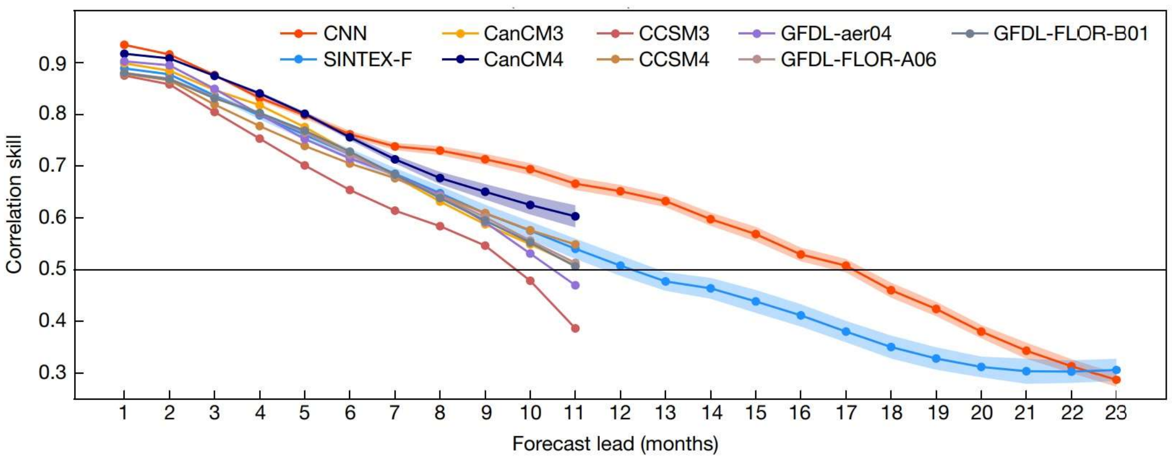

climate?

atus of machine learning at ECMWF

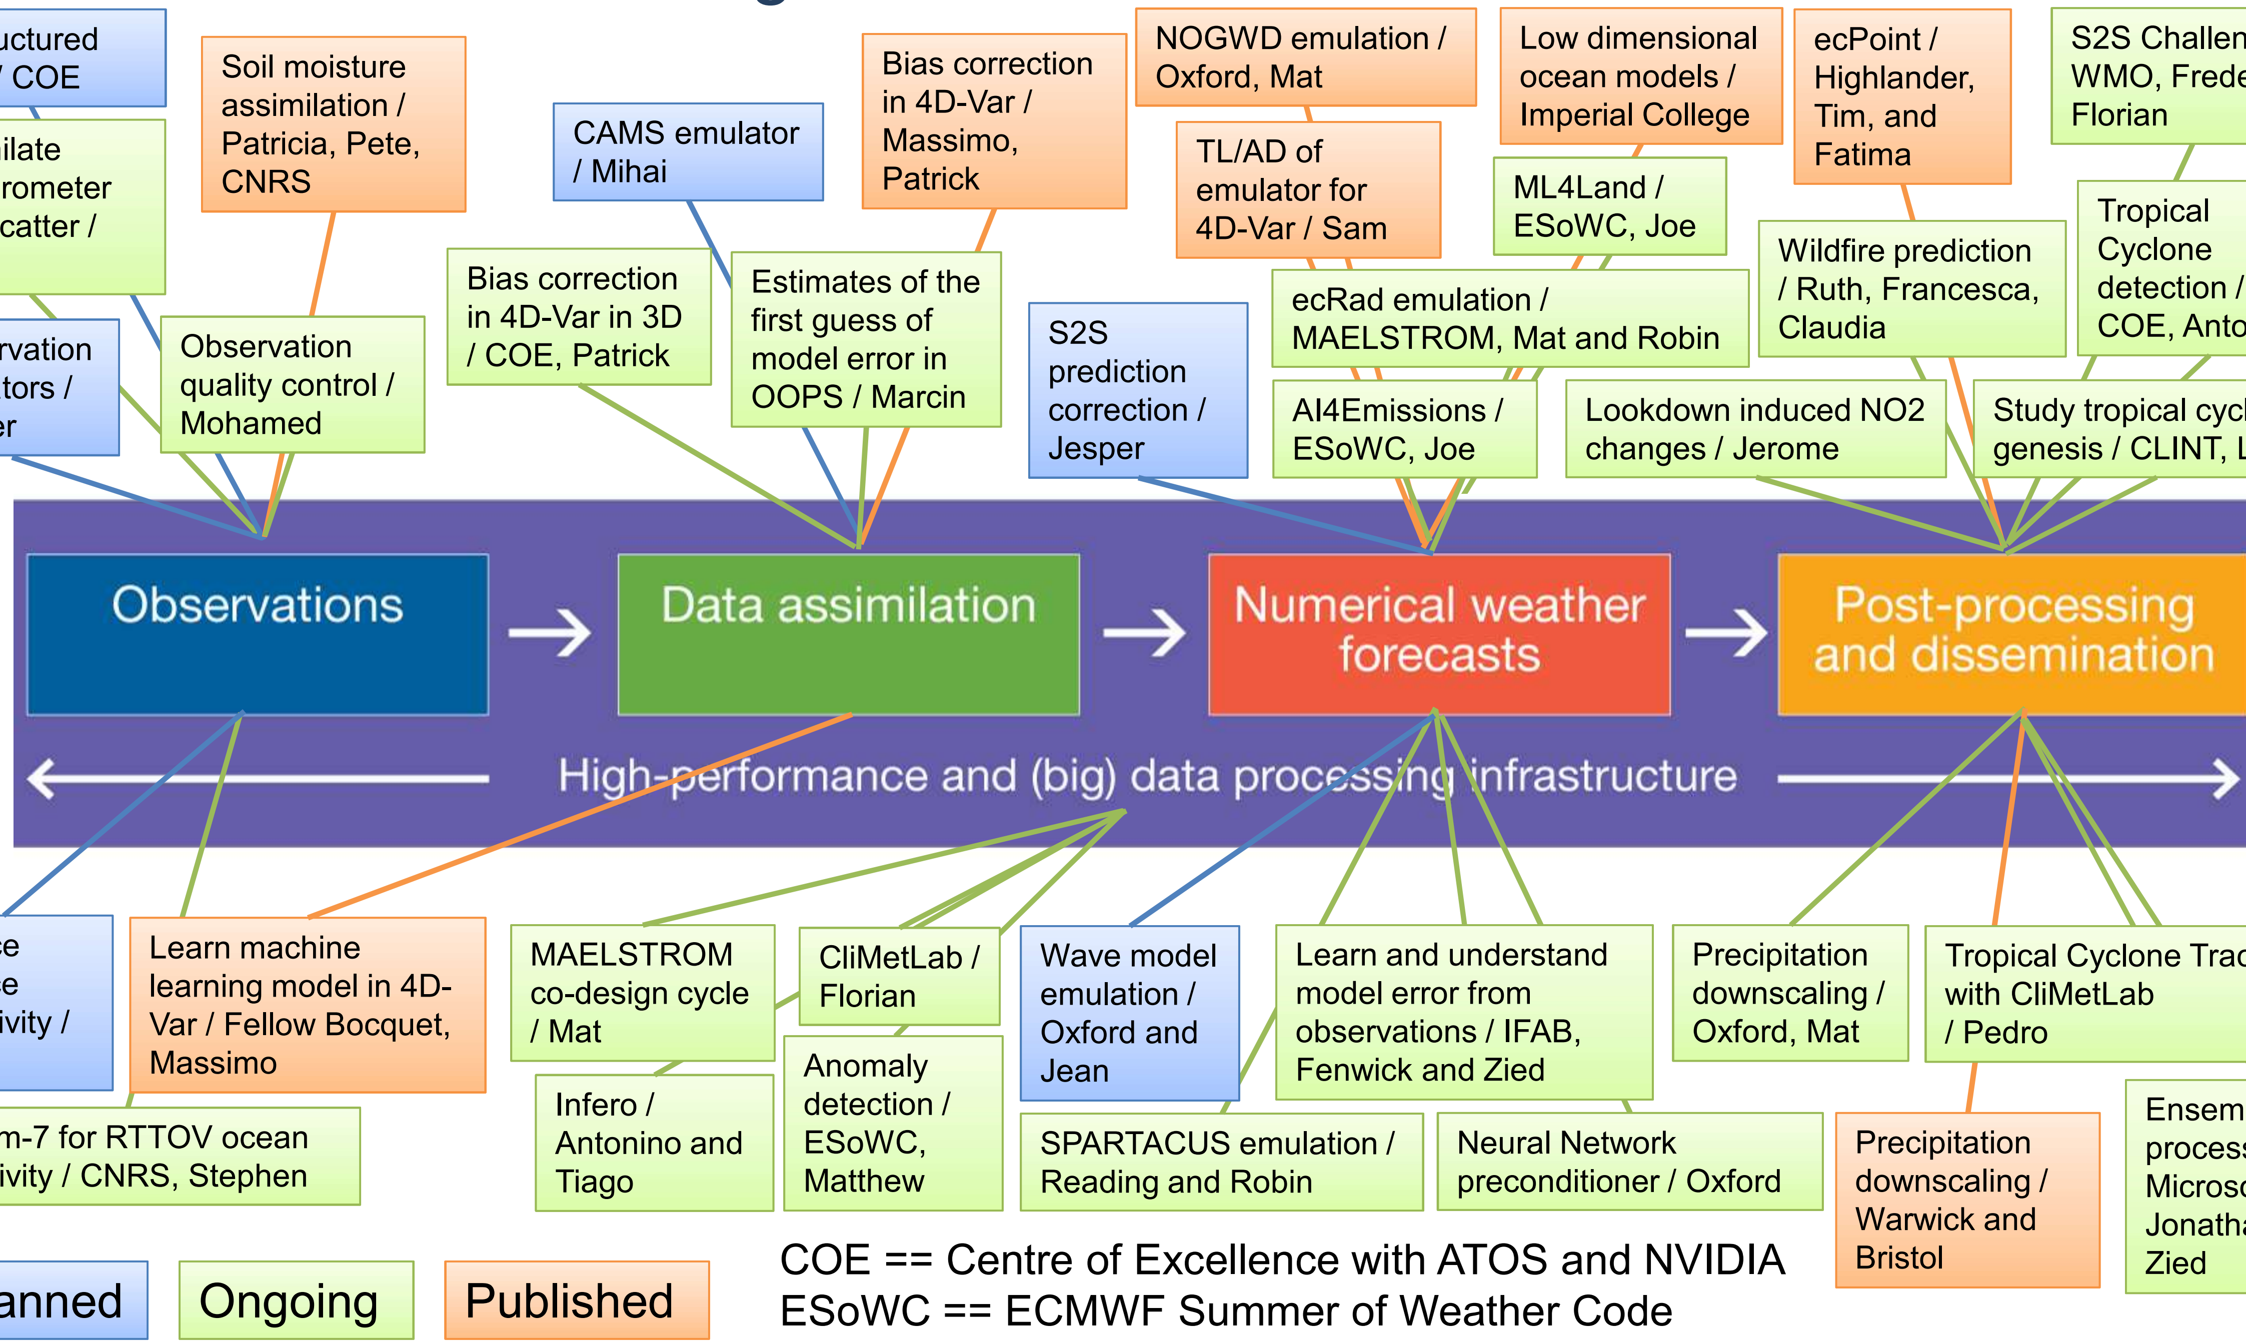

COE == Centre of Excellence with ATOS and NVIDIA
ESoWC == ECMWF Summer of Weather Code

ow bad is it to use machine learning in a changing climate?lopments

$$\frac{x}{t} = \sigma(y - x)$$

$$\frac{y}{t} = x(\rho - z) - y$$

$$\frac{z}{t} = xy - \beta z$$

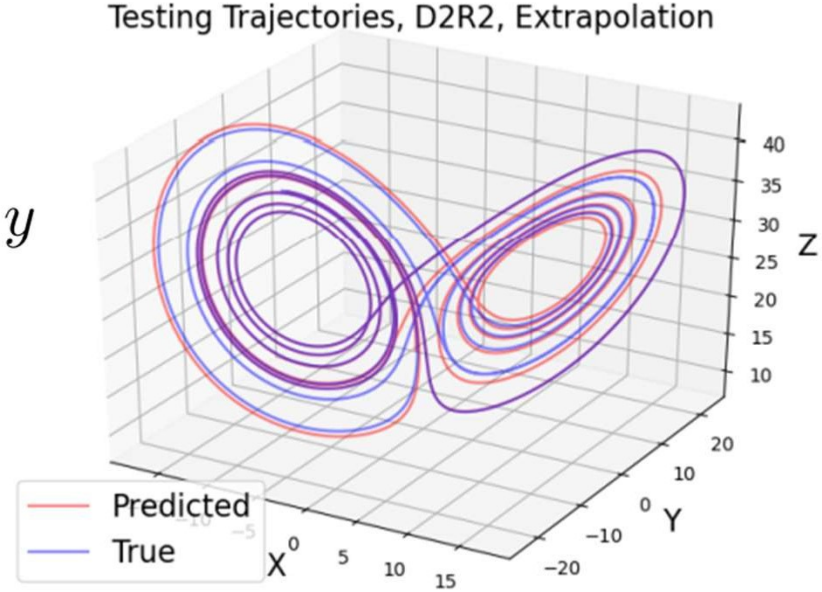

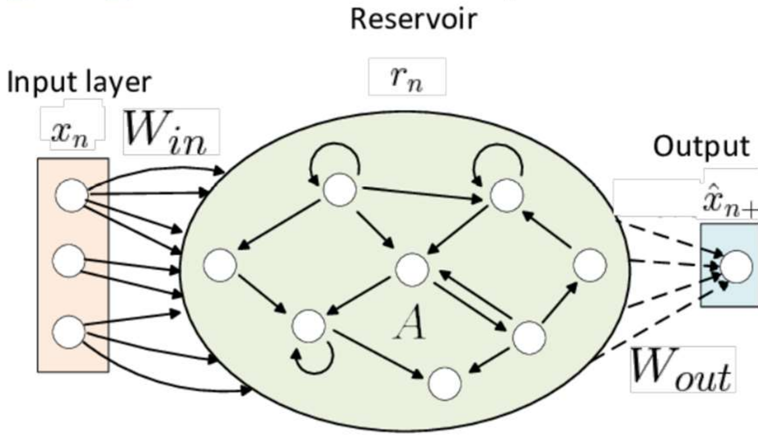

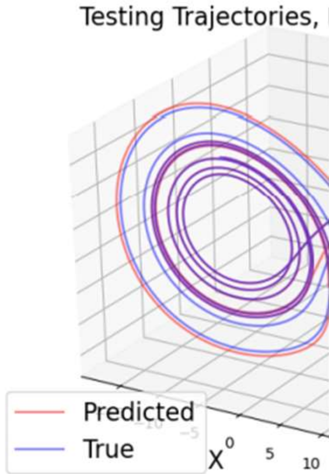

- et's train a machine learning tool in a changing climate
- et's start simple to be able to make clear statements → The Lorenz'63 model
- et's take two different approaches to learn the model from a truth trajectory:
  . Echo State Networks (Vlachas et al. 2020 and Chattopadhyay et al. 2020)
  . Domain-Driven Regularized Regression (D2R2; Pyle et al. 2021)
- et's assume that today's climate is the "left-lobe regime" and that climate change is kicking us into the 'two-lobe regime".
- Vhat if we only train from 1%, 2%, 5%… of the training data from the right lobe?

Pyle, Chantry, Palem, Palmer, Dueben, P

cience and tool developments

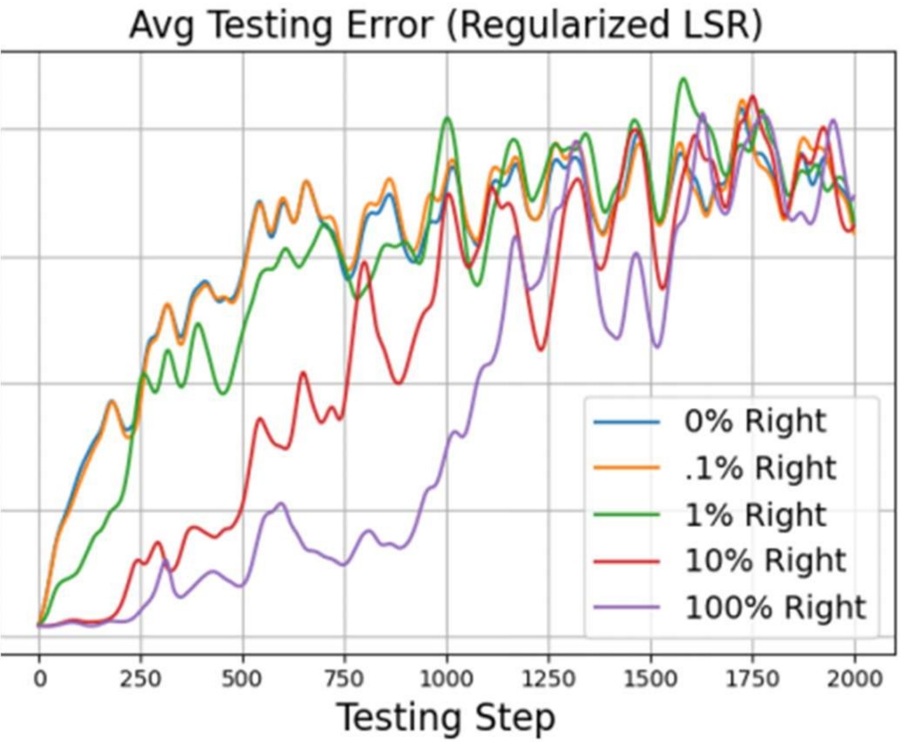

Echo State Network

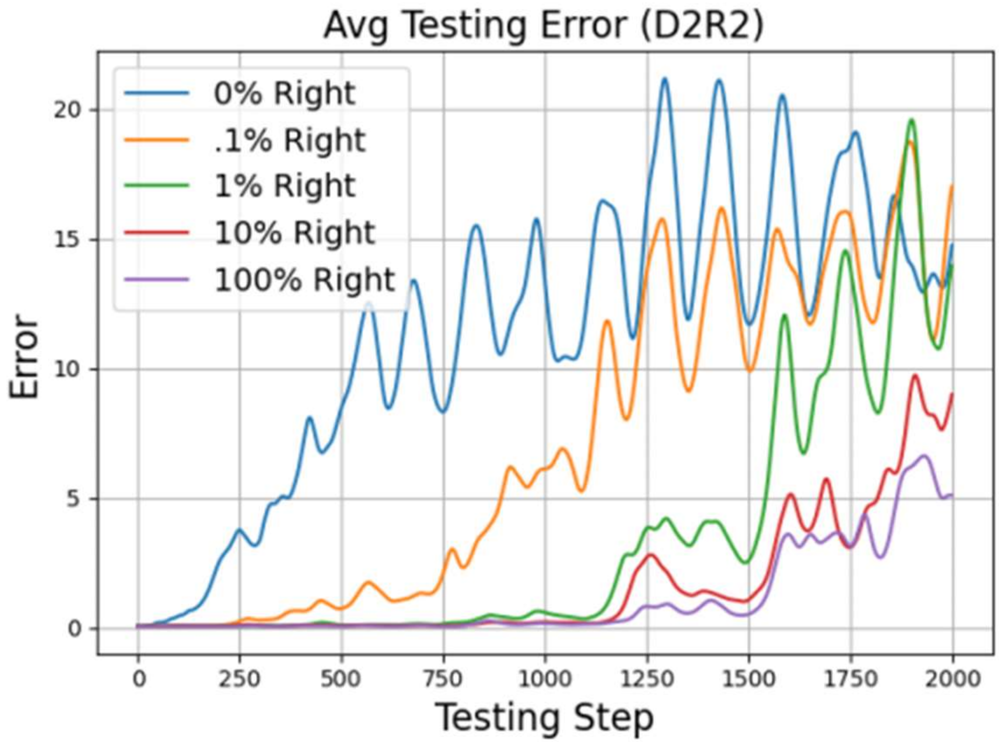

Regression Technique (D2R2)

The Echo State Network performs horrible unless you provide at least 10% of the data of the right lobe.
The regression technique needs a very small amount of the right lobe to perform well.

hysics informed machine learning, explainable AI and trustworthy AI need to be explored.

Pyle, Chantry, Palem, Palmer, Dueben, P

ow can you build trust in machine learning tools and make them iable?

stworthy AI, explainable AI and physics informed machine learning

re are several ways to incorporate physical knowledge into machine learning tools:

- Formulate the machine learning problem in a way that makes it physical (e.g. heating rate/fluxes for radi
- Change the architecture of the neural network
- Close the budget for the output variables or correct the outputs to fulfil the constraint
- Incorporate physical constraints into the loss function that is used for training

ere are also ways to evaluate whether the machine learning solution is reproducing the right phy

- Consider specific use cases and weather regimes
- Perform sensitivity tests on the inputs or outputs
- Test for physical reasoning (e.g. for extreme events)

chstein, M. et al. Deep learning and process understanding for data-driven Earth system science. Nature 566, 195–204 (

Govern, et al. Making the Black Box More Transparent: Understanding the Physical Implications of Machine Learning, Bu
e American Meteorological Society, 100(11), 2175-2199 (2019)

an we represent scale interactions with machine learning tools?

ather and climate modelling:
ls need to allow for scale interactions

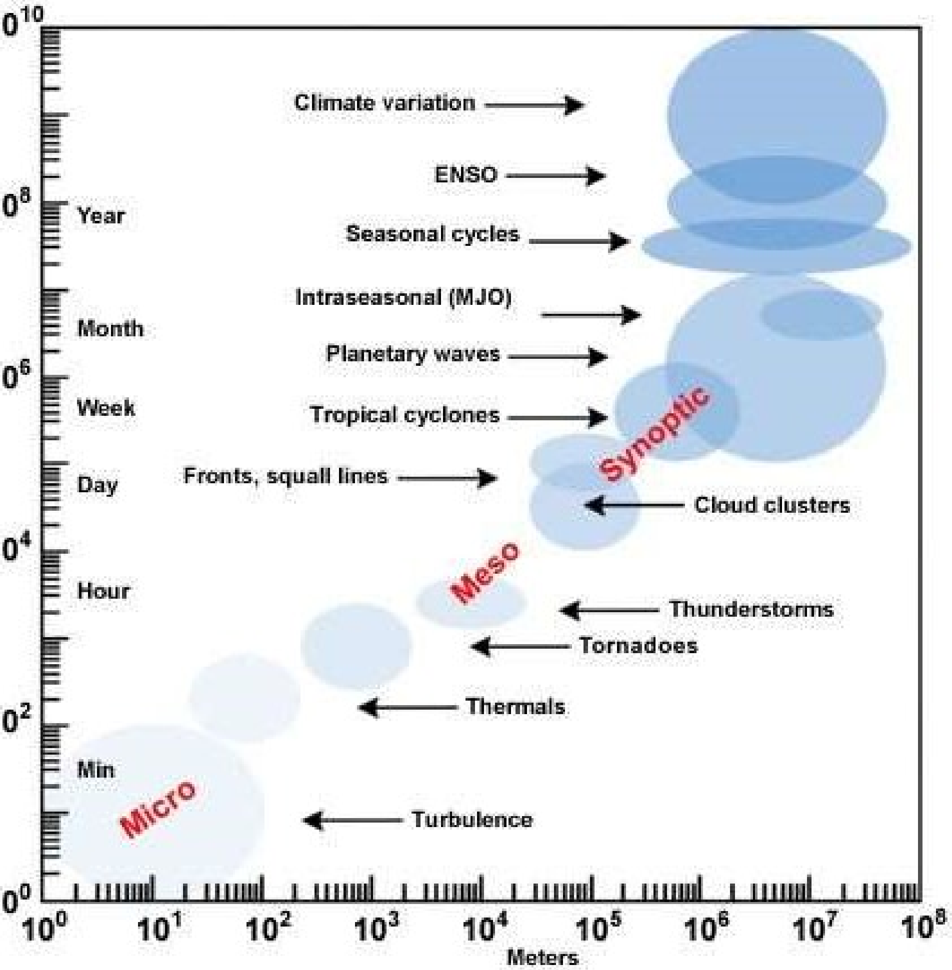

Source: UCAR

Machine learning:
Neural network tools allow for encoding/decoding stru

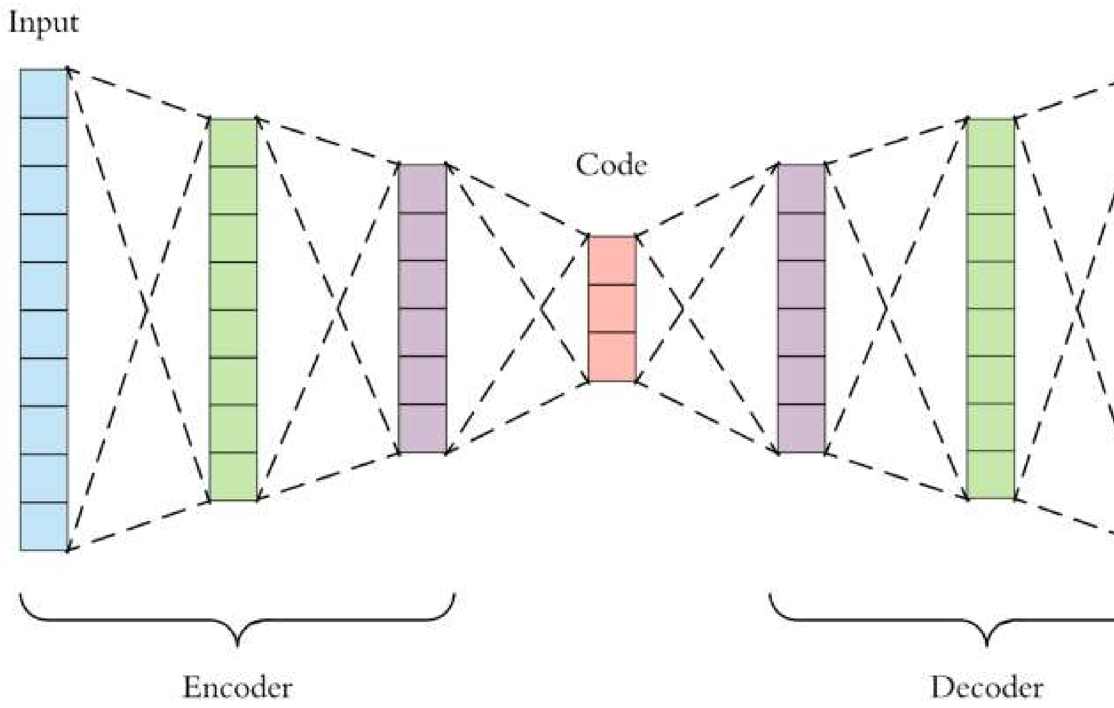

Source: https://towardsdatascience.com

n we use encoder/decoder networks to represent scale interactions?

recipitation down-scaling

oblem: Learn to map weather predictions from ERA5 reanalysis data at ~50 km resolution to E-OBS cal precipitation observations at ~10 km resolution over the UK.

se case: Eventually, apply the tool to climate predictions to understand changes of local precipitation ttern due to climate change.

ethod: Use Tru-NET with a mixture of ConvGru layers to represent spatial-temporal scale interactions d a novel Fused Temporal Cross Attention mechanism to improve time dependencies.

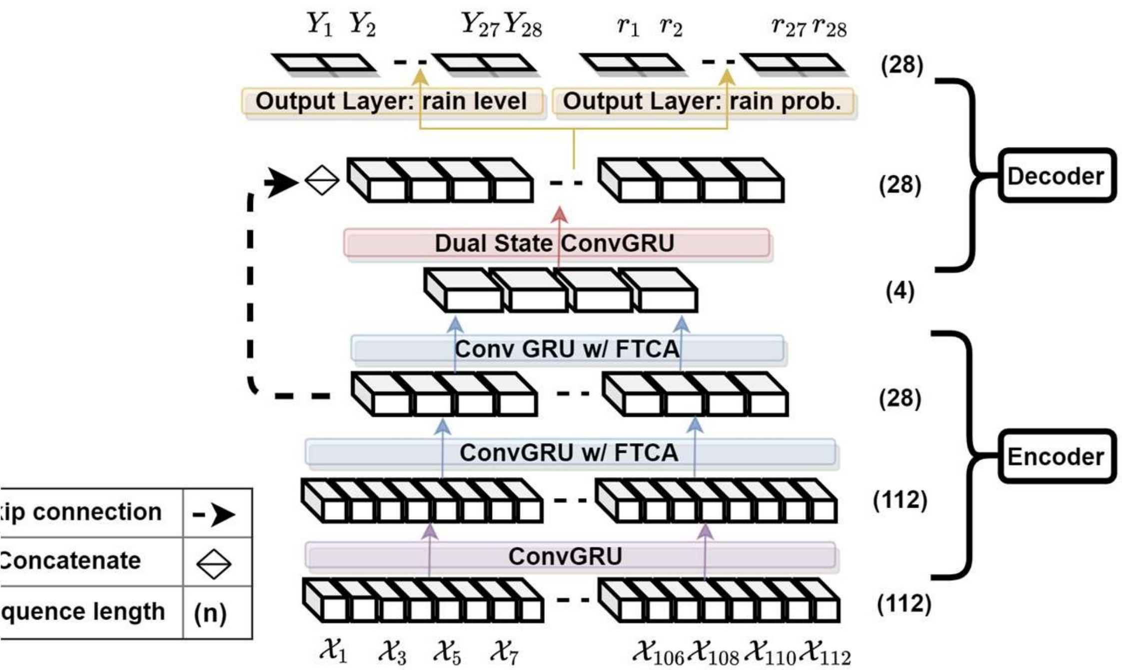

Model	RMSE
Conventional forecast model	3.627
Hierarchical Convolutional GRU	3.266
Tru-Net	3.081

Adewoyin, Dueben, Watson, He, Dutta http://arxiv.org/abs/200

ow to use machine learning?

It is often better to not replace the full system but rather to learn the "delta" of the most expensive or uncertain dynamics: No "all-in" but hybrid; no signal but delta
Watson, P. A. G.: Applying machine learning to improve simulations of a chaotic dynamical system using empirical error correction. Journal of Advances in Modeling Earth Systems, 11, 1402– 1417, 2019

It is often a good idea to learn the *error* since it is not "physical" and often measurable
Bonavita, M., & Laloyaux, P.: Machine learning for model error inference and correction. Journal of Advanc Modeling Earth Systems, 12, e2020MS002232, 2020

If you can learn the error, you can also learn the uncertainty representation
For example via dropout techniques, variational autoencoders or generative adversarial networks
Leinonen, J., Guillaume, A., & Yuan, T.: Reconstruction of cloud vertical structure with a generative advers network. Geophysical Research Letters, 46, 7035– 7044, 2019
See Hannah's talk

	Tripleclouds	SPARTACUS	Neural Network	Tripleclouds + Neural Network
tive	1.0	4.4	0.003	1.003

er, Hogan, Dueben, Mason https://arxiv.org/abs/2103.11919

PC efficiency and machine learning at scale

make efficient use of today's high performance computing hardware is tricky. Only a small number of ay's models can run on GPUs and most of the models run at <5% of the available peak performance.

ep learning tools are mostly based on dense linear algebra and reduced numerical precision.
DIA TensorCore on V100 GPUs perform matrix-matrix multiplications with:
8 TFlops for double precision
125 TFlops for half precision

e first machine learning application in weather and climate modelling has reached the exa-scale.
rsten Kurth et al.: Exascale deep learning for climate analytics. In Proceedings of the International Confer High Performance Computing, Networking, Storage, and Analysis (SC '18). IEEE Press, Article 51, 1–12, 20
rdon Bell Prize!

w much will we be able to learn when training from 1 petabyte of data using petascale supercomputin

Conclusions

There are a large number of application areas throughout the prediction workflow in weather ar climate modelling for which machine learning can make a difference.

The weather and climate community is still only at the beginning to explore the potential of machine learning (and in particular deep learning) at scale and there are challenges to be face

However, an approach that combines collaborations, meetings, scientific studies, targeted projects, shared datasets, software and hardware developments should allow us to overcome most of the challenges in the medium-term future.

lease do not forget to register for the ESA-ECMWF Workshop on Machine Learning for arth System Observation and Prediction – 15-18 November – https://www.ml4esop.esa.int

any thanks! **Peter.Dueben@ecmwf.int** **@PDueber**

ECMWF | EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS

The strength of a common goal

achine learning in three communities

v did the view on machine learning change from 2018 until today?

bold Machine Learning scientist:
achine learning will replace everything"
"Machine learning will replace everything, look here…"

HPC hardware developer:
achine learning will dominate future HPC developments"
"Here is our new machine learning hardware, please use it"

sceptical weather and climate domain scientist:
achine learning is just a wave going through…"
"Machine learning is just a method…"

there is still more that can be done with customised machine learning tools that are easy to use at scale

hallenges for machine learning in weather and climate modelling

fferent sets of tools for domain (Fortran on CPUs) and machine learning scientists (Python on GPUs)
Training and tool development (e.g. CliMetLab)

f-the-shelf machine learning tools are often not sufficient for weather and climate applications
Science, benchmark datasets and tool developments

aining datasets are often not good enough while the data size is huge
Benchmark datasets

e still need to learn how to scale up to petascale supercomputers to make the most of machine learnin
Projects such as MAELSTROM and benchmark datasets

tegration of machine learning tools into the conventional numerical weather prediction workflow is dif
Science and tool developments (e.g. Infero)

achine learning tools need to be updated in model cycles
Science (e.g. Transfer Learning)

achine learning tools need to be reliable (extrapolating?) for use in operational predictions
Science (e.g. explainable AI, trustworthy AI or physics-informed networks)

n we use deep learning hardware for conventional models?

ew operational model configuration:

del configuration	Relative Cost
uble precision 91 levels	100%
gle precision 91 levels	57.9%
uble precision 137 levels	155.5%
gle precision 137 levels	87.5%

gle precision is used for operational dictions at ECMWF since May 2021

change from double to single precision and m 91 to 137 vertical levels allows to reduce ts *and* improve predictions

Tropical cyclone intensity (core pressure) bias
Red: Single precision and 137 vertical levels
Blue: Double precision and 91 vertical levels

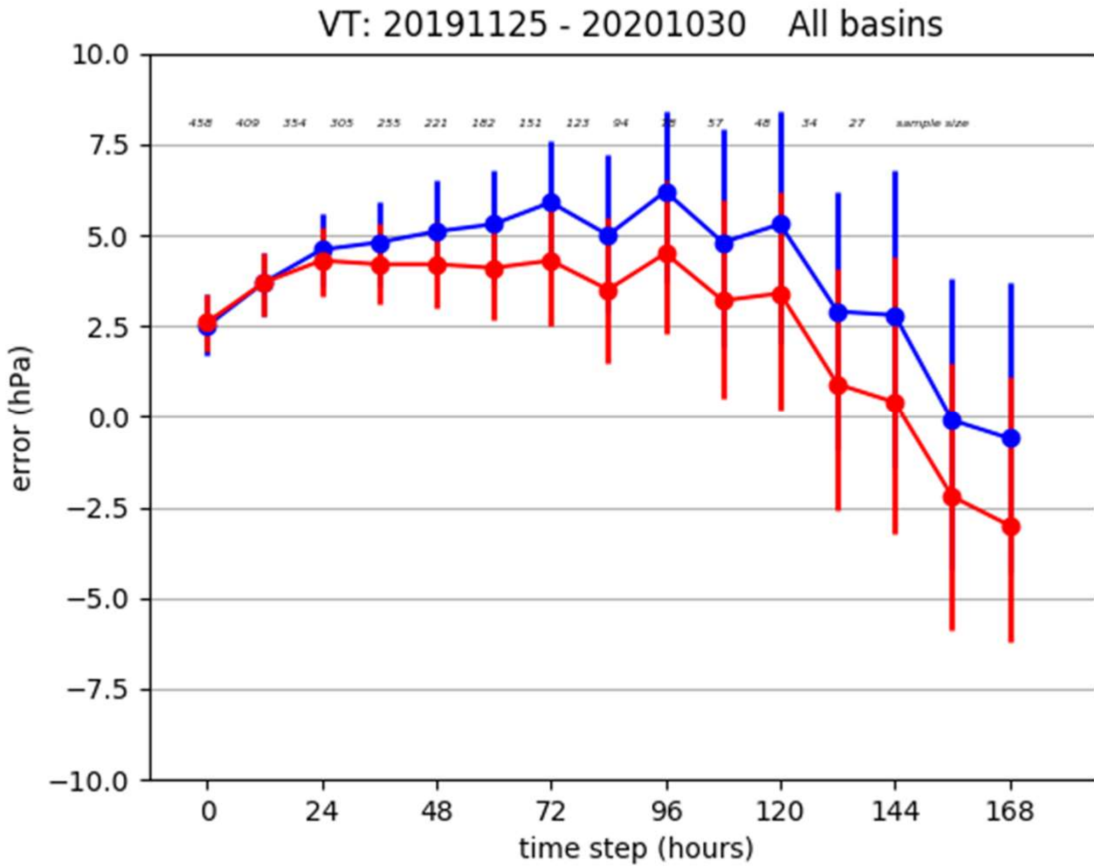

Dueben and Palmer 2014 → Lang et al. submitted to

an we use deep learning hardware for conventional models?

- Machine learning accelerators are focussing on low numerical precision and high floprats.
- Example: TensorCores on NVIDIA Volta GPUs are optimised for half-precision matrix-matrix calculations with single precision output.
  → 7.8 TFlops for double precision vs. 125 TFlops for half precision

Can we use TensorCores within our models?

Relative cost for model components for a non-hydrostatic model at 1.45 km resolution:

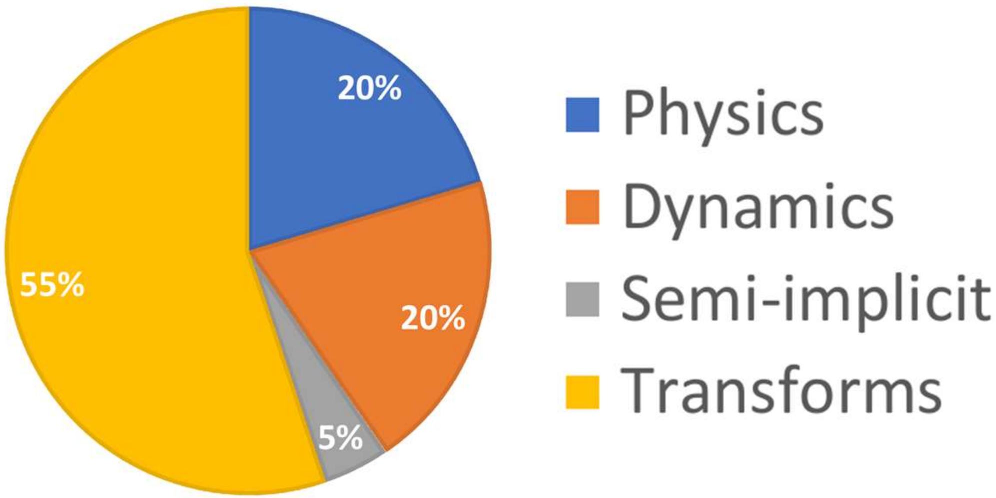

- The Legendre transform is the most expensive kernel. It consists of a large number of standard matrix-matrix multiplications.
- If we can re-scale the input and output fields, we can use half precision arithmetic.

alf precision Legendre Transformations

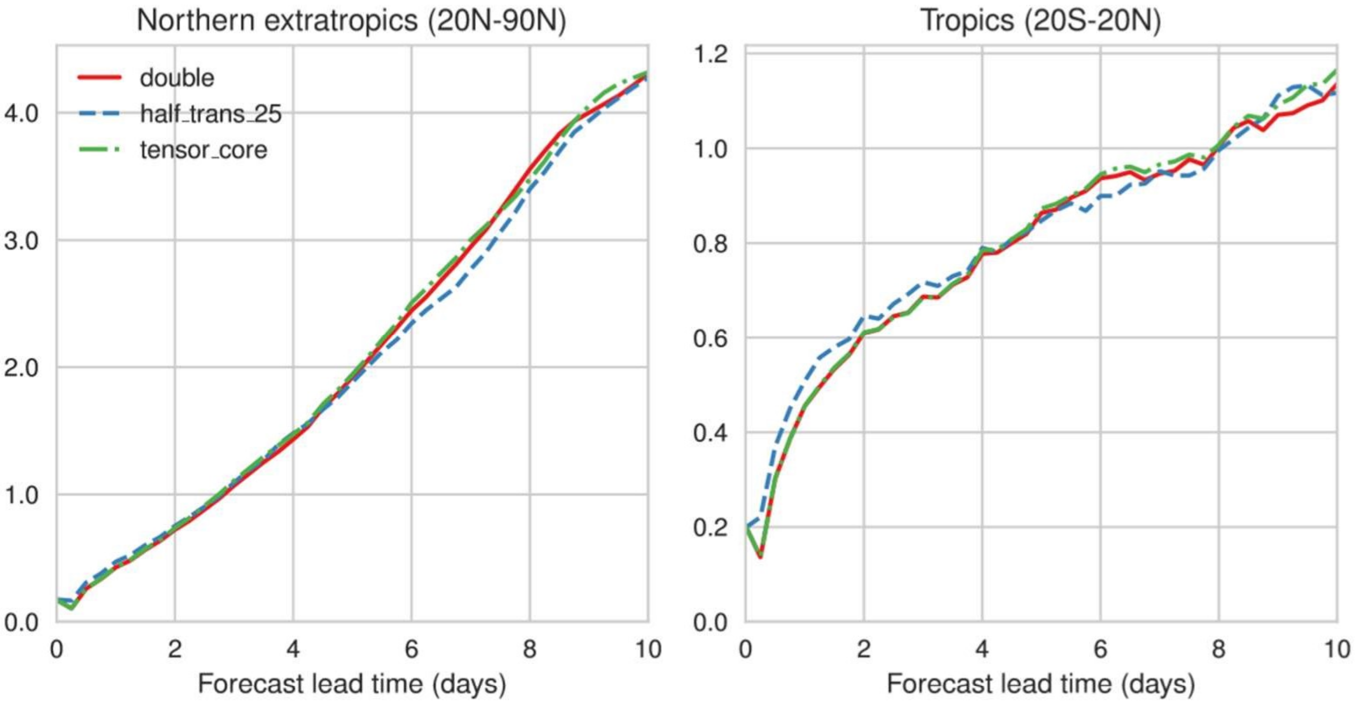

ot-mean-square error for geopotential height at 500 hPa at
m resolution averaged over multiple start dates. *Hatfield,*
antry, Dueben, Palmer Best Paper Award PASC2019

e simulations are using an emulator to reduce precision
awson and Dueben GMD 2017) and more thorough
gnostics are needed.

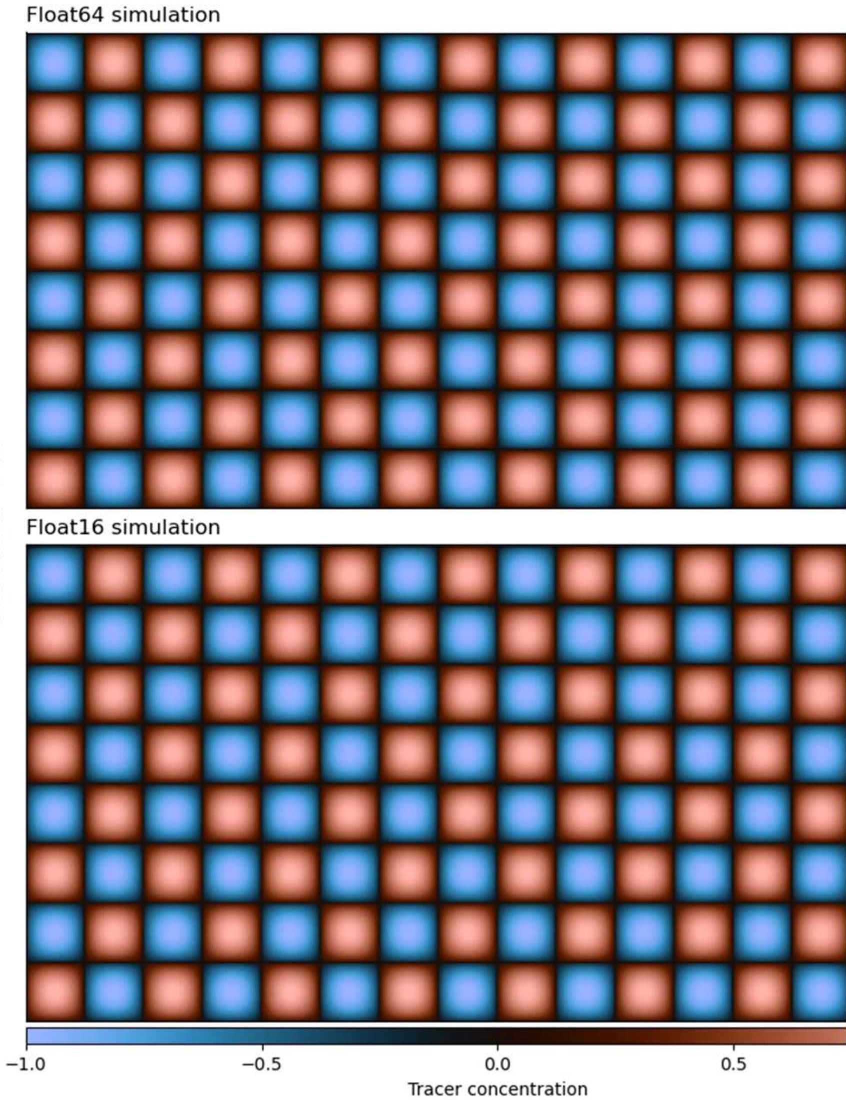

Results from Sam Hatfield on Fuga
(many thanks to Hirofumi Tomita!)
and from Milan Kloewer on Isamba

hallenges for machine learning in weather and climate modelling

erent sets of tools for domain (Fortran on CPUs) and machine learning scientists (Python on GPUs)

the-shelf machine learning tools are often not sufficient for weather and climate applications

ning datasets are often not good enough while the data size is huge

still need to learn how to scale up to petascale supercomputers to make the most of machine learning

gration of machine learning tools into the conventional numerical weather prediction workflow is diff

chine learning tools need to be updated in model cycles

chine learning tools need to be reliable (extrapolating?) for use in operational predictions

achine learning for bias correction

- During data-assimilation the model trajectory is “synchronised” with observations
- Model error can be diagnosed when comparing the model with (trustworthy) observations

proach: Learn model error for a given model state using machine learning

efit: Correct for model error and understand model deficiencies

estion: What happens when the model is upgraded and the error pattern change?

ution: More work on transfer learning needs to be done

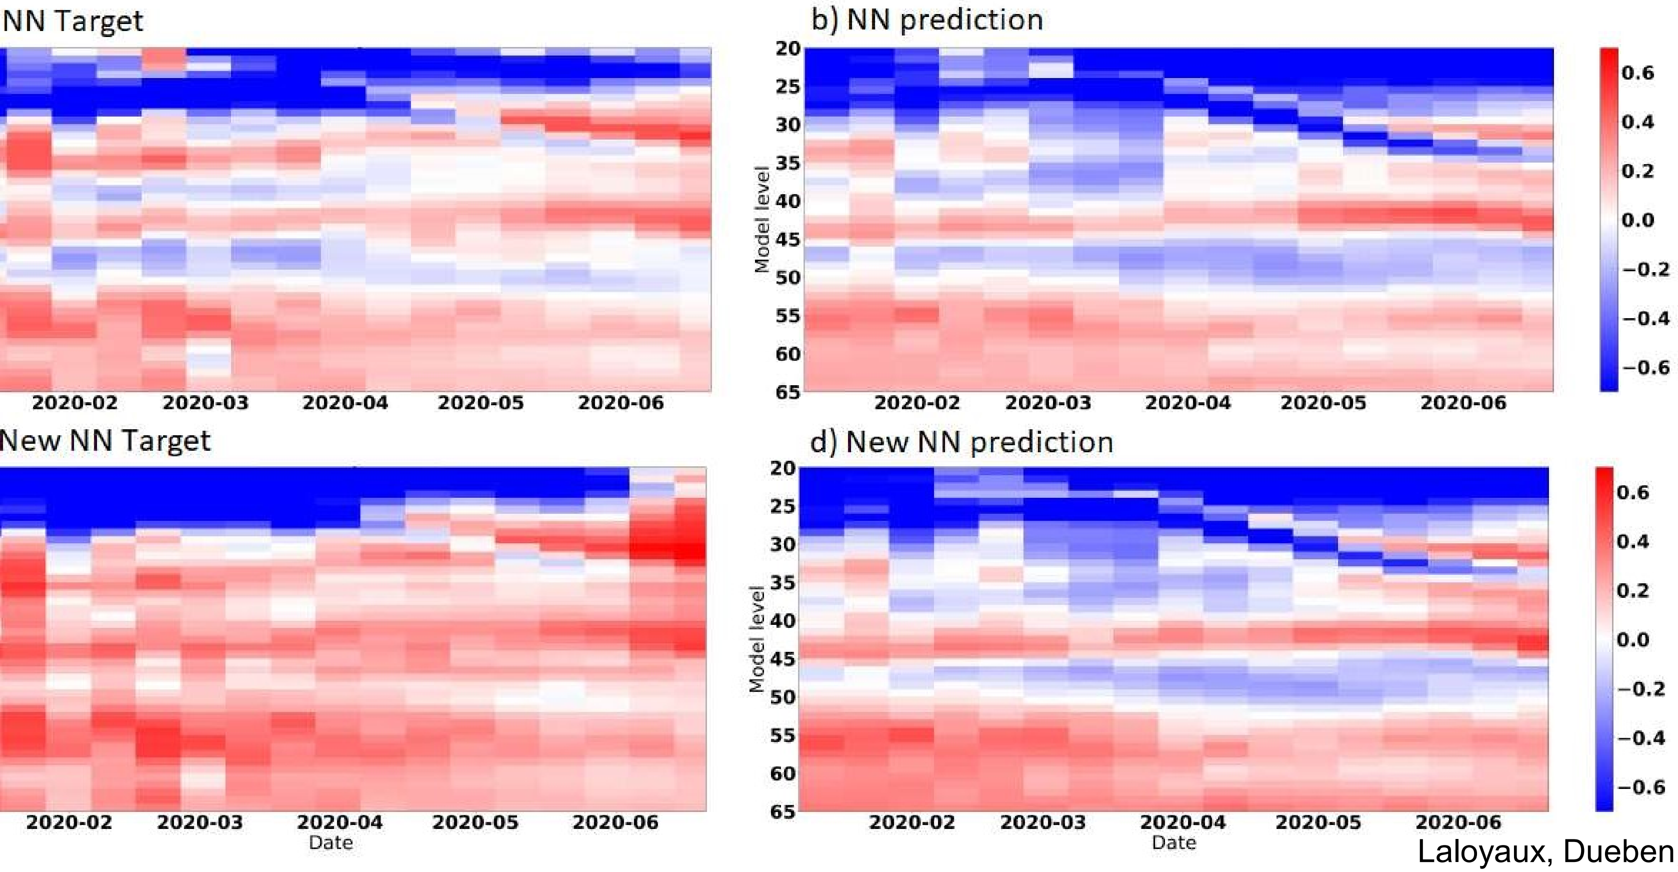

Laloyaux, Dueben, Bonavita @ ECMWF + Kurth and Hall @

emulate parametrisation schemes

thod:
- Store input/output data pairs of a parametrisation scheme
- Use this data to train a neural network
- Replace the parametrisation scheme by the neural network within the model

y would you do this?
ural networks are likely to be much more efficient and portable to erogenous hardware

tive area of research:
evallier et al. JAM 1998, Krasnopolsky et al. MWR 2005, Rasp et al. PNAS 2018, nowitz and Bretherton GRL 2018…

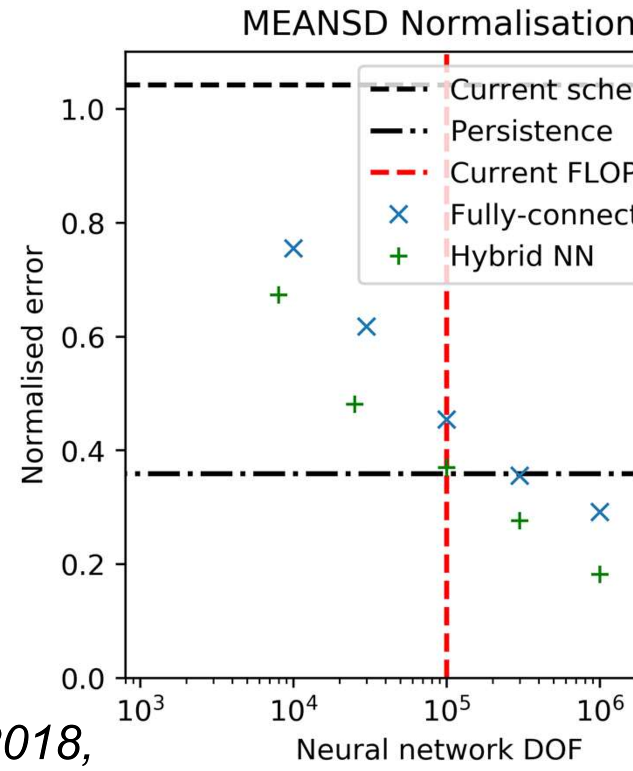

emulate the non-orographic gravity wave drag within the Integrated Forecasting System (IFS)
antry, Hatfield, Dueben, Polichtchouk and Palmer https://arxiv.org/abs/2101.08195

sults:
- Nice relationship between neural network complexity and error reduction
- Similar cost when used within IFS on CPU hardware and 10 times faster when used offline on GPUs
- Emulator was used successfully to generate tangent linear and adjoint code within 4D-Var data assimilatio
  Hatfield, Chantry, Dueben, Lopez, Geer, Palmer in preparation
- Forecast error can be reduced when training with more angles and wavespeed elements

precondition the linear solver

Linear solvers are important to build efficient semi-implicit time-stepping schemes for atmosphere and ocean mo
However, the solvers are expensive.
The solver efficiency depends critically on the preconditioner that is approximating the inverse of a large matrix.

**n we use machine learning for preconditioning, predict the inverse of the matrix and reduce the number o
rations that are required for the solver?**

stbed: A global shallow water model at 5 degree resolution but with real-world topography.
thod: Neural networks that are trained from the model state and the tendencies of full timesteps.

preconditioner:

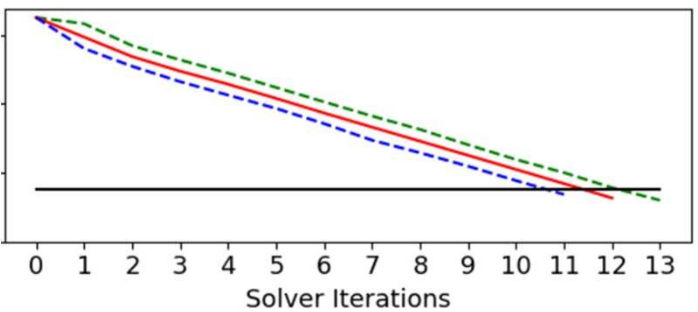

Machine learning preconditioner:

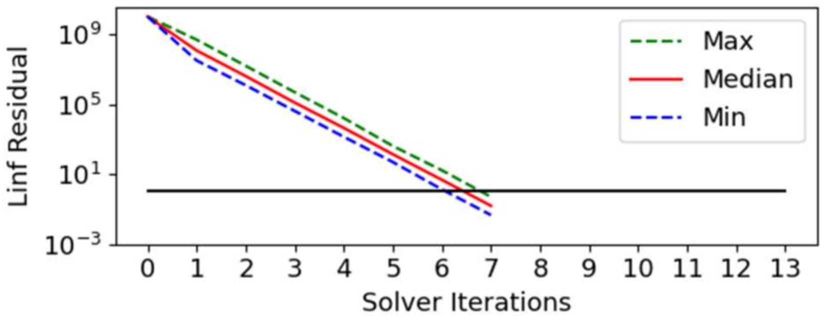

Implicit Richardson precondi

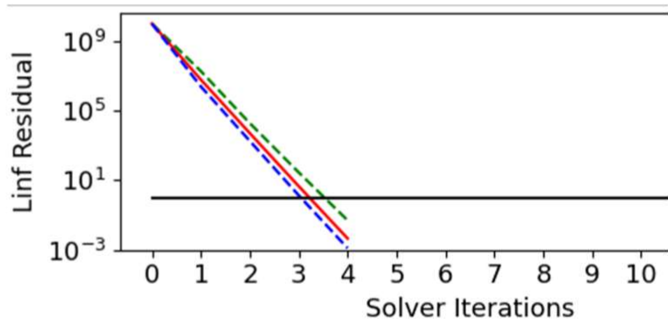

**urns out that the approach (1) is working and cheap, (2) interpretable and (3) easy to implement
en if no preconditioner is present.**

ost-processing and dissemination: ***ecPoint*** to post-process rainfall predictic

- Jse forecast data as inputs
- Train against worldwide rainfall observations
- mprove local rainfall predictions by accounting or sub-grid variability and weather-dependent biases
- Jse decision trees as machine learning tool

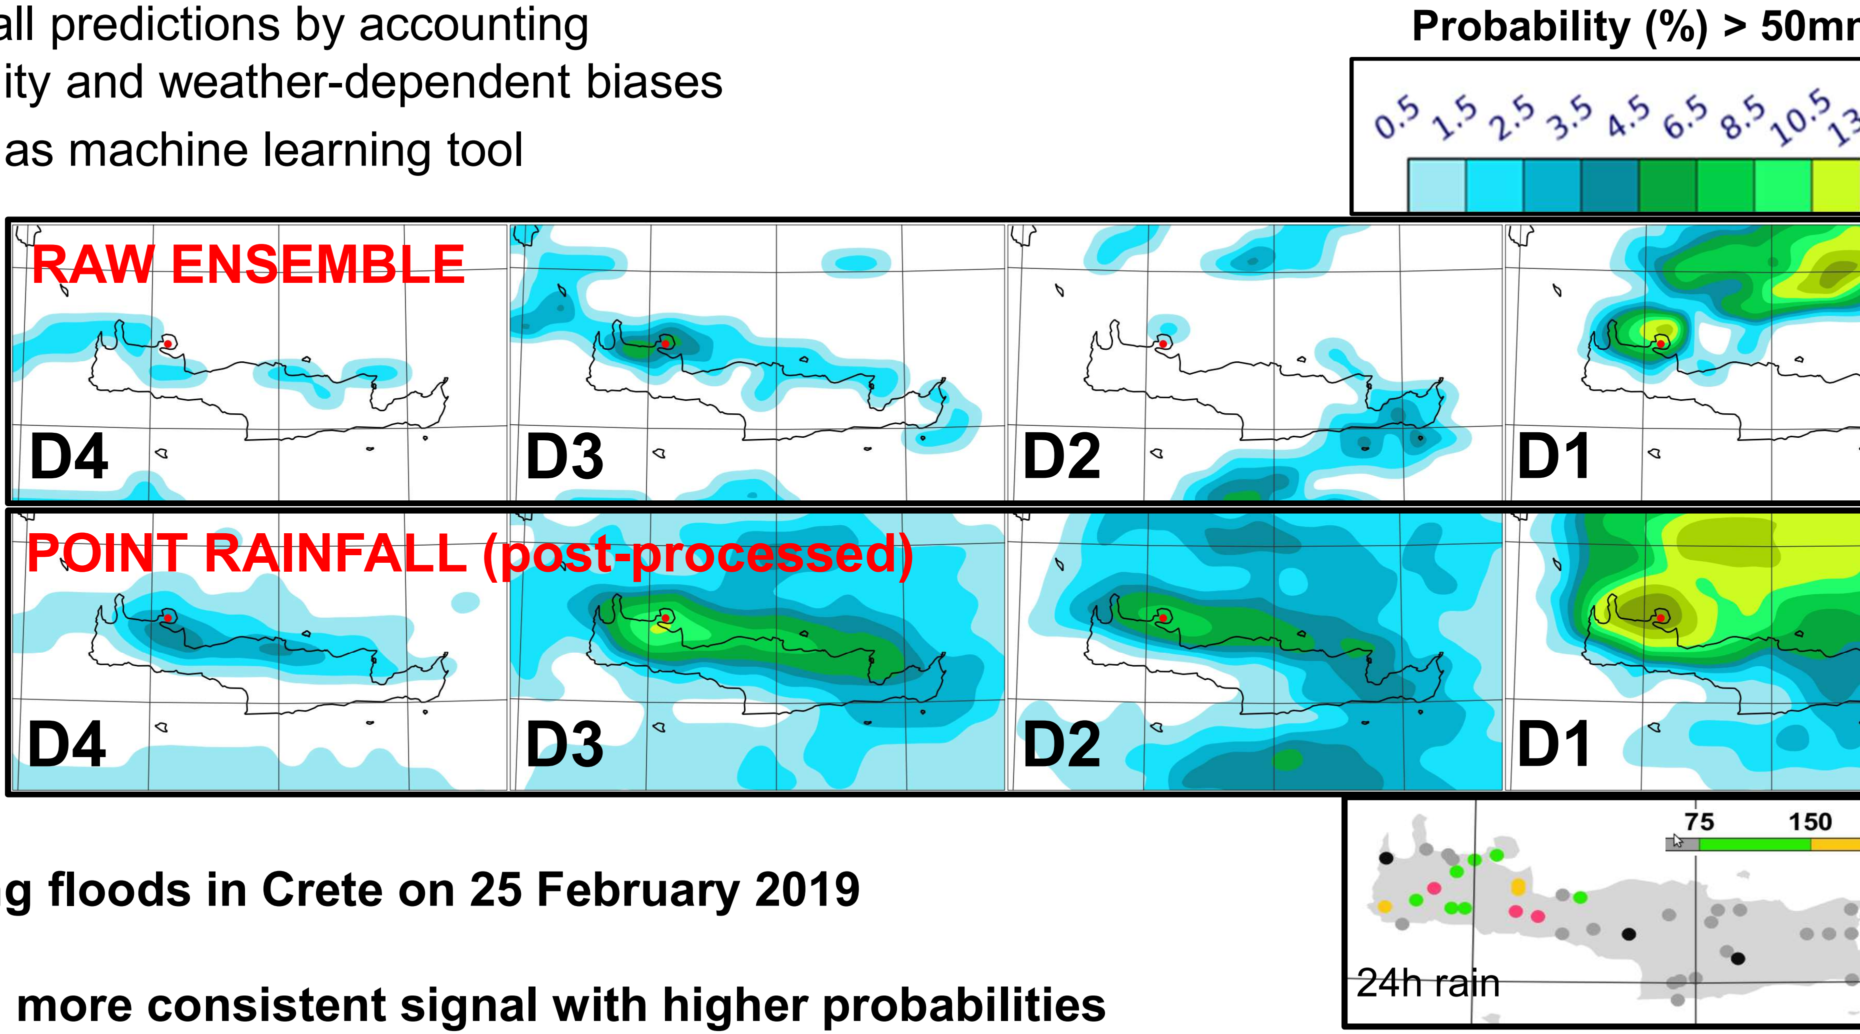

mple: Devastating floods in Crete on 25 February 2019

nefits: Earlier and more consistent signal with higher probabilities

Timothy Hewson and Fatim

post-processing ensemble predictions

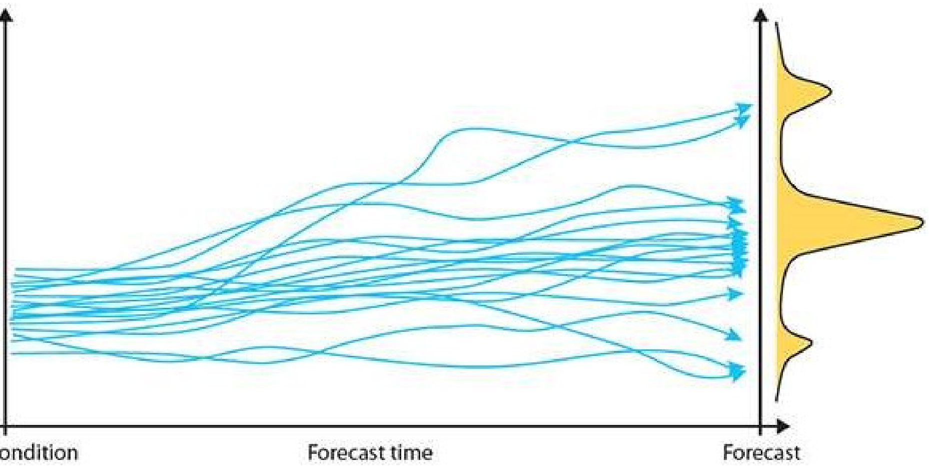

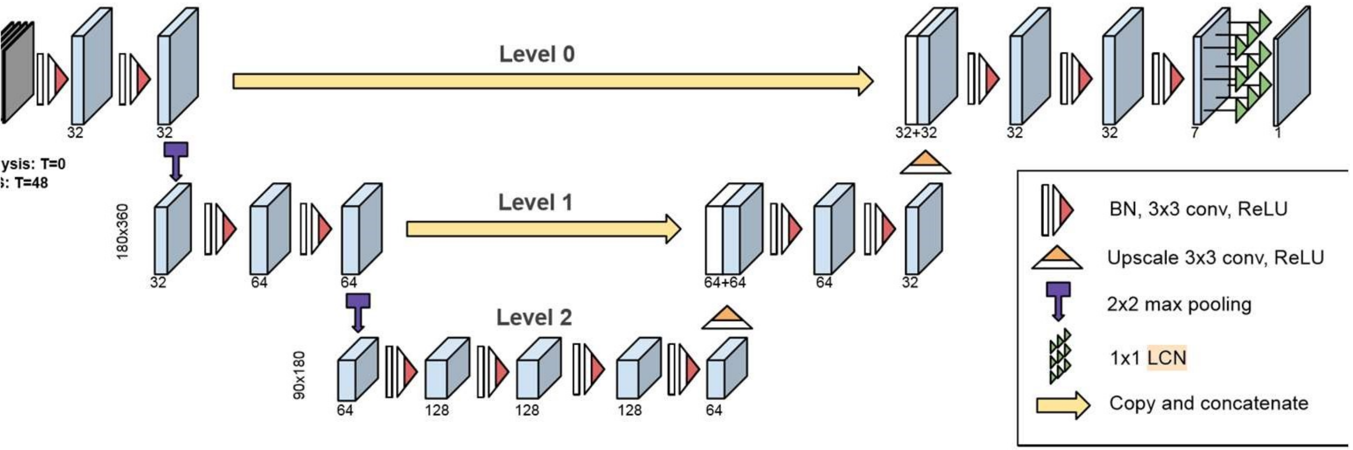

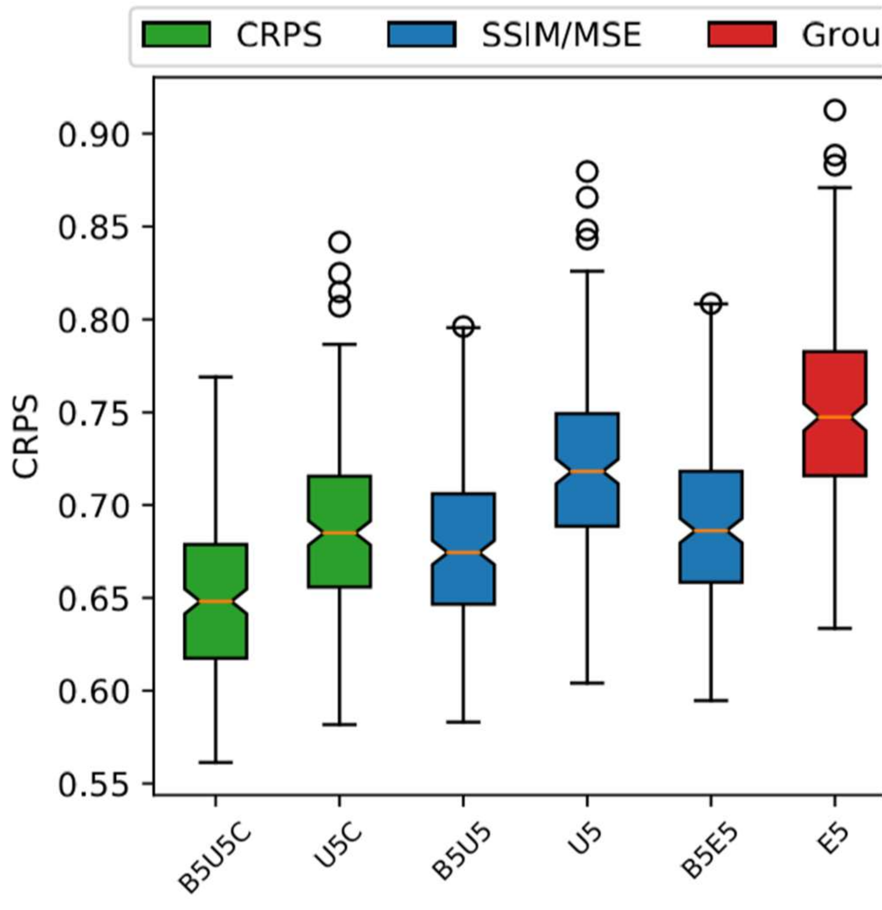

(a) T850

semble predictions are important but expensive.

n we improve ensemble skill scores from a small number of ensemble members via deep learning?

Use global fields of five ensemble members as inputs.

Correct the ensemble scores of temperature at 850 hPa and Z500 hPa for a 2-day forecast towards a full 1
member ensemble forecast.

obabilistic down-scaling

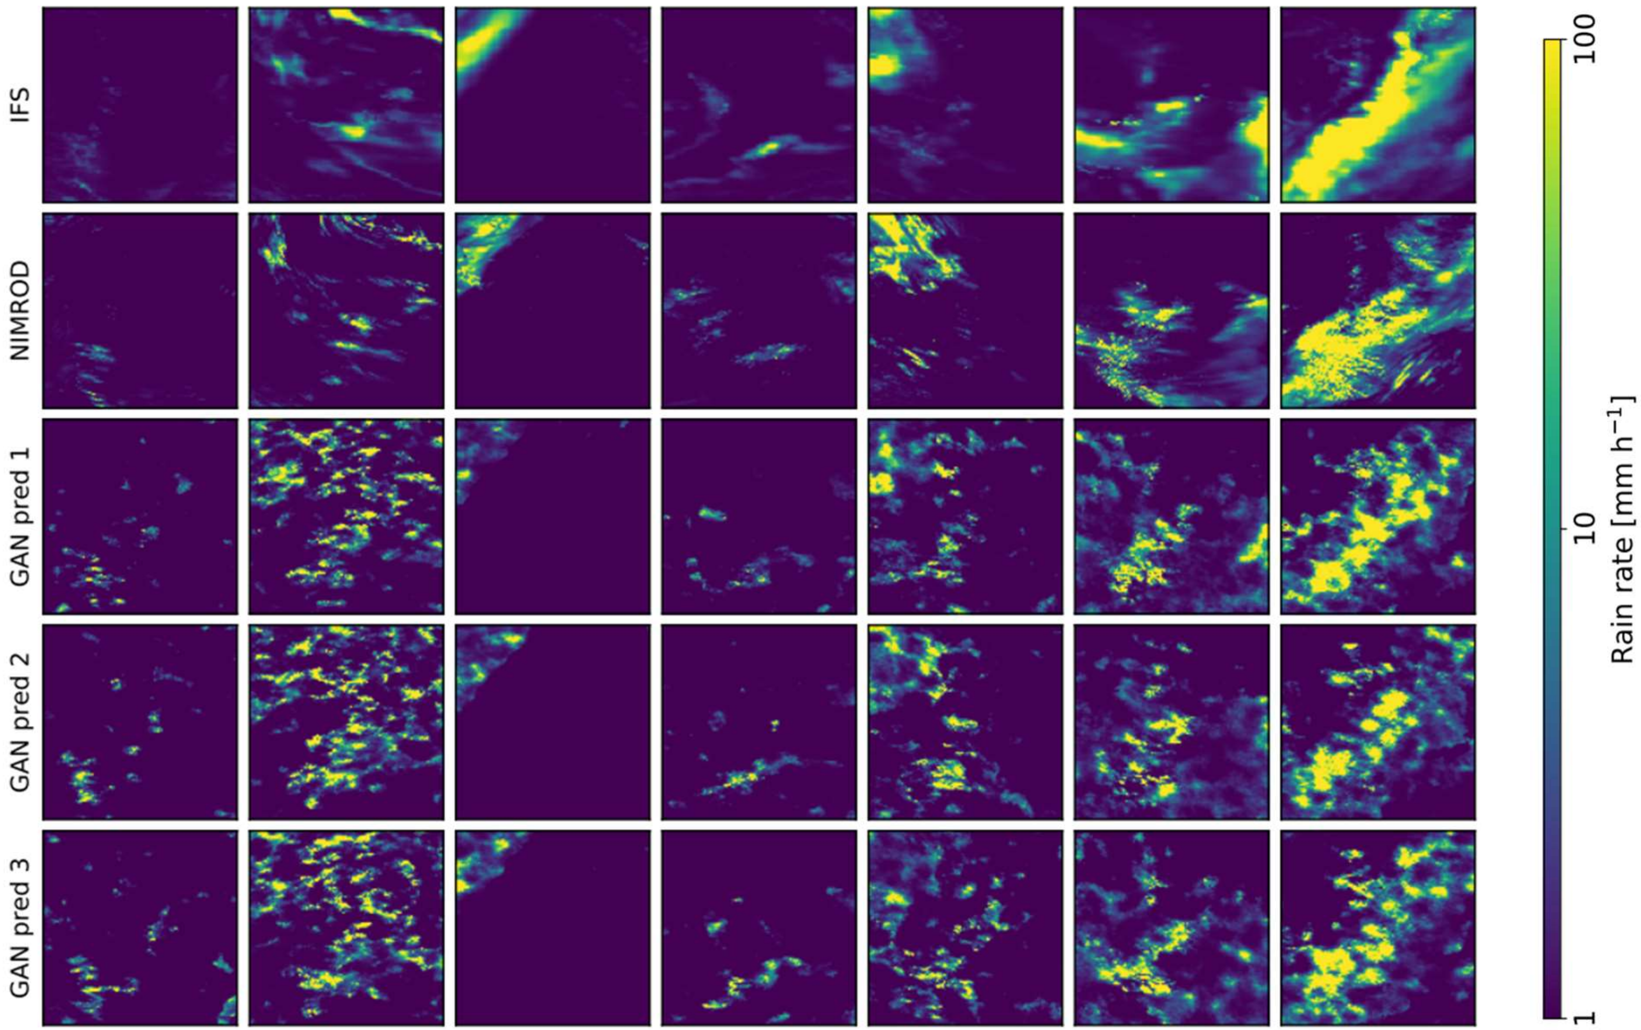

Map IFS model data at ~10 km resolution to NIMROD precipitation observations at ~1 km resolution
Test Generative Adversarial Networks (GANs) and Variational Autoencoders (VAs)
Generate ensembles to represent the uncertainty of the mapping.

hallenges for machine learning in weather and climate modelling

fferent sets of tools for domain (Fortran on CPUs) and machine learning scientists (Python on GPUs)
Machine learning roadmap via training and tool development (e.g. CliMetLab)

f-the-shelf machine learning tools are often not sufficient for weather and climate applications
Machine learning roadmap, MAELSTROM and COE via benchmark datasets and tool developments

aining datasets are often not good enough while the data size is huge
MAELSTROM via benchmark datasets

e still need to learn how to scale up to petascale supercomputers to make the most of machine learnin
MAELSTROM via co-design cycle

tegration of machine learning tools into the conventional numerical weather prediction workflow is dif
Science and tool developments, COE, and tool development (e.g. Infero)

achine learning tools need to be updated in model cycles
Science and tool developments and COE via Transfer Learning

achine learning tools need to be reliable (extrapolating?) for use in operational predictions
Science and tool developments