The next 10 years of machine learning at ECMWF

--

An introduction to the machine learning roadmap

Peter Dueben

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

ECMWF | EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS

The strength of a common goal

Machine learning: What and why?

What is machine learning?

- Computer algorithms that improve automatically through learning from data without being explicitly programmed
- Learn non-linear mappings between fields (supervised)
- Extract information from data (unsupervised)

Why now?

- Increase in data volume and knowledge
- New computing hardware
- New machine learning software

Why weather and climate?

- Complex, non-linear, many components
- A lot of data
- Supercomputing application

Why do we need a roadmap?

- Only at the beginning and challenges ahead
- Infrastructure needs
- Many applications
  → launch with the Strategy

Slide from Torsten Hoefler (ETH)

Many application areas for machine learning across ECMWF

Hydrology models and identification of human influence

Gravity wave drag emulation

ecPoint precipitation post-processing

Automated quality control of analysis

Clustering and categorisation for predictability and model errors

Hybrid ML/conventional land-surface models

Pollution anomaly detection and pollution downscaling

Post-processing for extreme predictions in Copernicus

Anomaly detection

Radiation emulation

ML preconditioning

Learn forecasts from observations

Feature detection

Observations → Data assimilation → Numerical weather forecasts → Post-processing and dissemination

← High-performance and (big) data processing infrastructure →

ML preconditioning for 4D-Var

Bias correction and model learning

Optimise data access

CliMetLab

Precipitation downscaling

Use land surface observations and data assimilation

Bias correction and gap filling for aerosol observations

Anomaly detection and workflow for IT infrastructure

Fuse observations and predictions for sea ice products

Mapping of non-Gaussian to Gaussian distributions for data assimilation

Optimise chillers for HPC

Ensemble post-processing

However, our roadmap does not provide an outline for specific scientific projects…

Plenty of activities at ECMWF

1. Centre of Excellence in Weather & Climate Modelling with ATOS and ECMWF supported by AMD, Mellanox, NVIDIA and DDN
2. European Weather Cloud together with EUMETSAT
3. Opening of the data archive
4. New projects starting (MAELSTROM, AI4Copernicus and CLINT)
5. CliMetLab to simplify access to climate and meteorological data
6. Two workshops organised in 2019 and 2020
7. Machine learning seminar series in 2020
8. AI and Machine Learning coordinator
9. Destination Earth

…

MAELSTROM

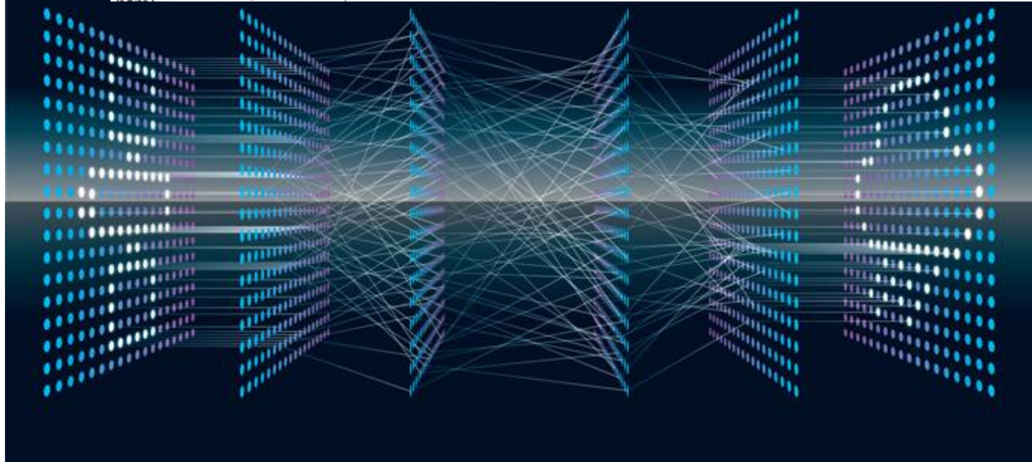

High-level objectives:

Objective 1
Explore machine learning applications across the weather and climate prediction workflow and apply them to improve model efficiency and prediction quality.

Objective 2
Expand software and hardware infrastructure for machine learning.

Objective 3
Foster collaborations between domain and machine learning experts with the vision of merging the two communities.

Objective 4
Develop customised machine learning solutions for Earth system sciences that can be applied to various applications and at scale on current and future supercomputing infrastructure.

Objective 5
Train staff and Member and Co-operating State users and organise scientific meetings and workshops.

Specific milestones:

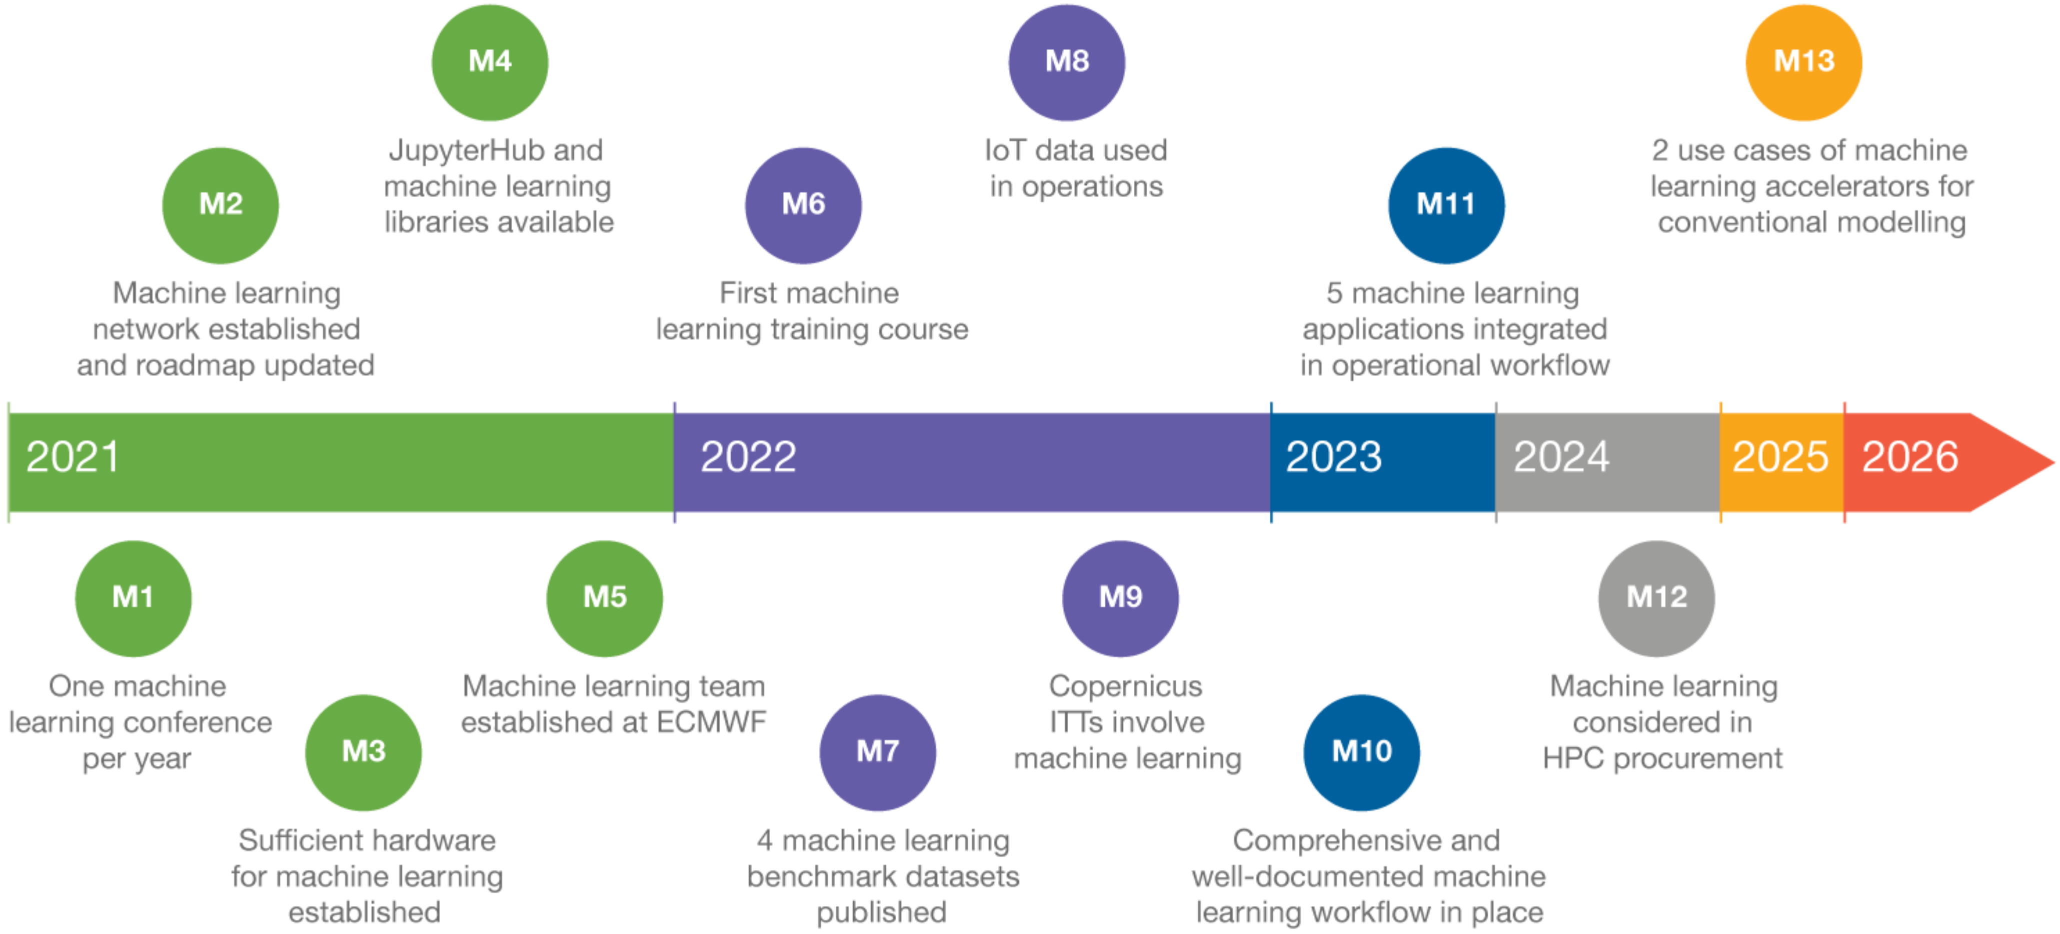

Vision 2031

- It is difficult to distinguish between machine learning and domain sciences
- Data handling fully capable to serve machine learning needs
- Fully supported diagnostic tools via trustworthy AI
- Physical constraints can be represented in deep learning
- Use of machine learning as easy and normal as data re-gridding
- Unsupervised learning and causal discovery used on a regular basis
- Machine learning solutions from end-users integrated in workflow

Challenges and milestones

Different philosophy for domain and machine learning scientists
Approach: *Support close collaborations // study explainable AI, trustworthy AI and physics informed machine learning*

For many applications off-the-shelf machine learning tools will not be sufficient
Approach: *Foster cross-disciplinary collaborations // develop customised machine learning tools // Benchmark Datasets*

Difficult to learn from observations and to *improve* models
Approach: *Learn from and exploit data assimilation // learn boundary conditions from observations*

Data avalanche
Approach: *Anticipate data access and channelise requests // efficient use of heterogeneous hardware*

Different set of tools (e.g. Fortran on CPUs vs. Python on GPUs)
Approach: *Training // Software // Hardware*

Integrate machine learning tools into the conventional NWP and climate service workflow
Approach: *Centralised tools and efforts // embed efforts into the scalability project*

- M1 — One machine learning conference per year
- M2 — Machine learning network established and roadmap updated
- M3 — Sufficient hardware for machine learning established
- M4 — JupyterHub and machine learning libraries available
- M5 — Machine learning team established at ECMWF
- M6 — First machine learning training course
- M7 — 4 machine learning benchmark datasets published
- M8 — IoT data used in operations
- M9 — Copernicus ITTs involve machine learning
- M10 — Comprehensive and well-documented machine learning workflow in place
- M11 — 5 machine learning applications integrated in operational workflow
- M12 — Machine learning considered in HPC procurement
- M13 — 2 use cases of machine learning accelerators for conventional modelling

2021 | 2022 | 2023 | 2024 | 2025 | 2026

Specific milestones:

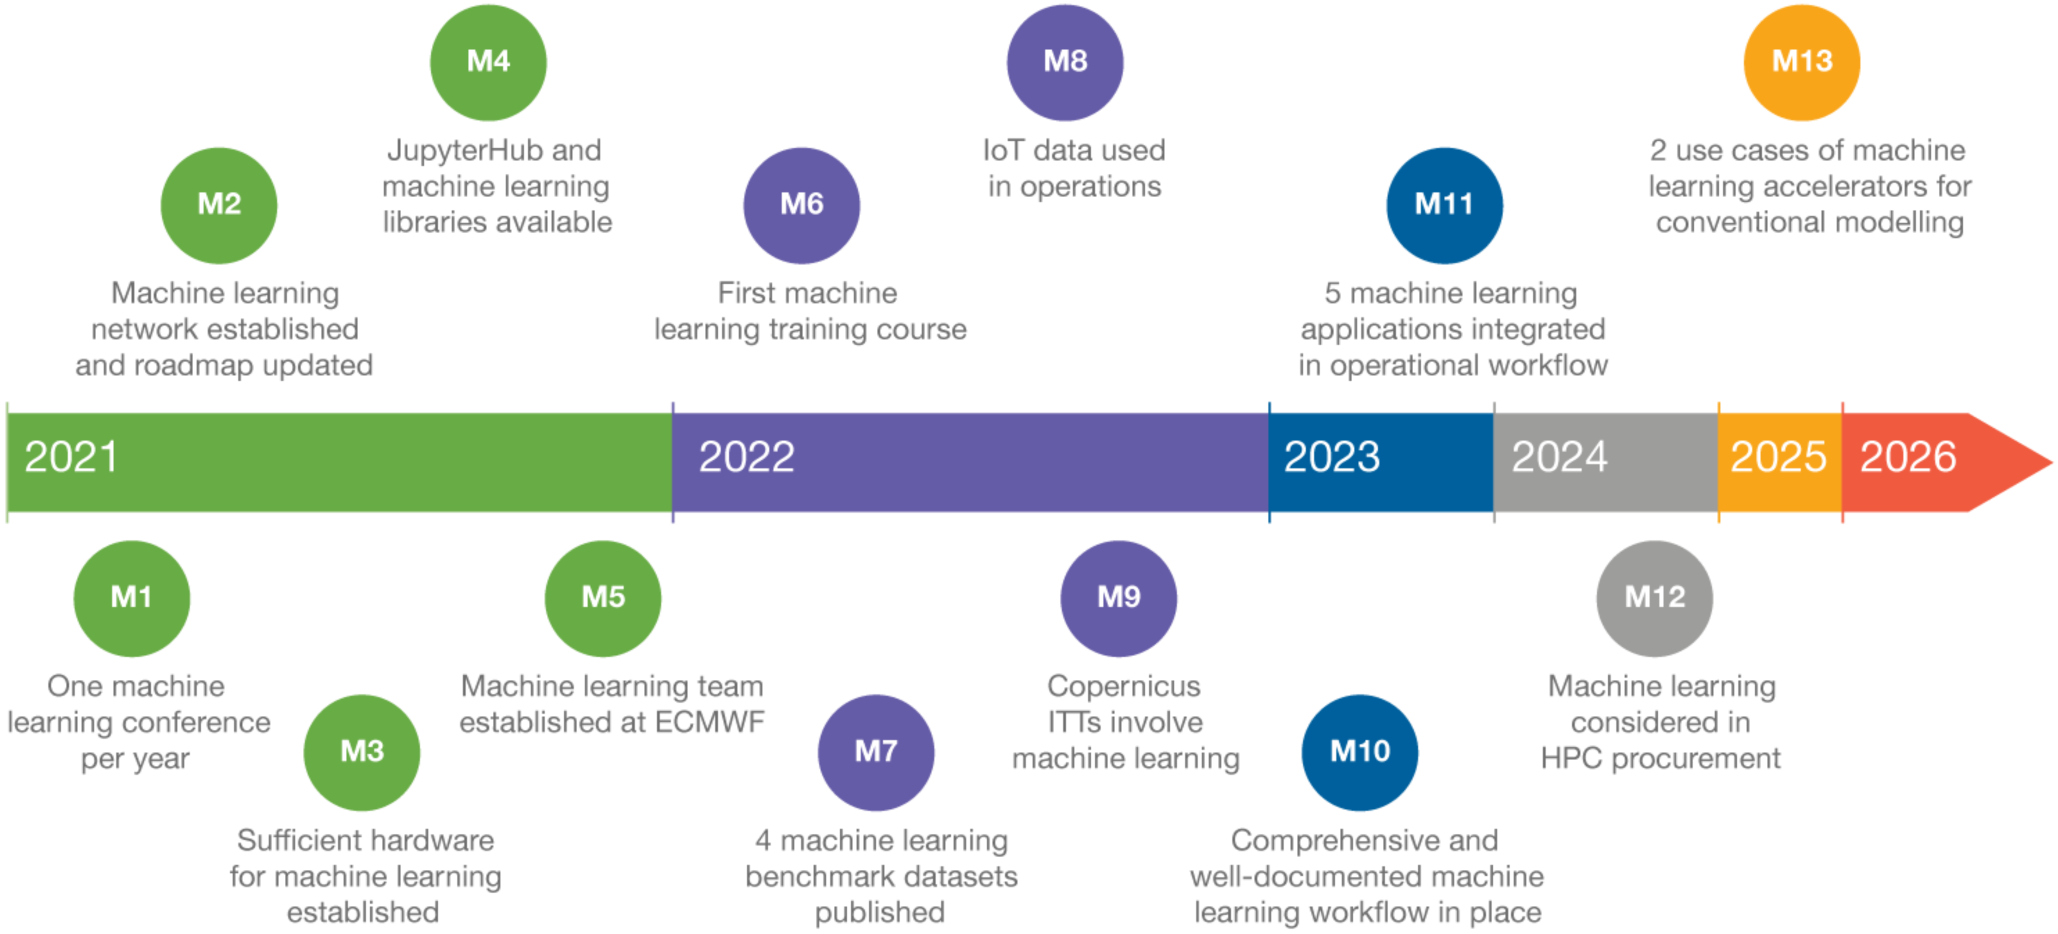

Vision 2031

- It is difficult to distinguish between machine learning and domain sciences
- Data handling fully capable to serve machine learning needs
- Fully supported diagnostic tools via trustworthy AI
- Physical constraints can be represented in deep learning
- Use of machine learning as easy and normal as data re-gridding
- Unsupervised learning and causal discovery used on a regular basis
- Machine learning solutions from end-users integrated in workflow

To emulate parametrisation schemes

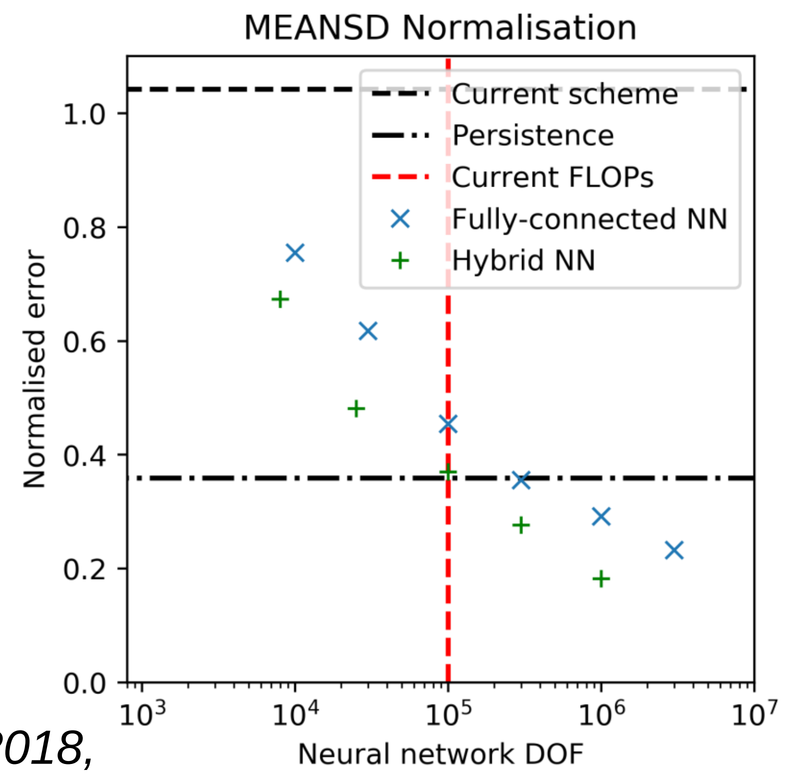

Method:

- 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

Why would you do this?

Neural networks are likely to be much more efficient and portable to heterogenous hardware

Active area of research:

Chevallier et al. JAM 1998, Krasnopolsky et al. MWR 2005, Rasp et al. PNAS 2018, Brenowitz and Bretherton GRL 2018…

We emulate the non-orographic gravity wave drag within the Integrated Forecasting System (IFS)

Chantry, Hatfield, Dueben, Polichtchouk and Palmer https://arxiv.org/abs/2101.08195

Results:

- 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 assimilation
  Hatfield, Chantry, Dueben, Lopez, Geer, Palmer in preparation
- Forecast error can be reduced when training with more angles and wavespeed elements

To emulate parametrisation schemes

Implemented workflow with great efforts from Mat Chantry and Sam Hatfield:

1. Output generated within IFS based on GRIB format
2. Initial training on CPUs
3. Guinea pig project on European Weather Cloud
4. Long-lasting search for the ideal network architecture
5. Dense neural networks are implemented in IFS in Fortran Code
6. Speed-up is limited as the code blocking is not ideal

Targeted workflow in 2023:

1. Output generated within IFS based on GRIB format
2. Output retrieval and IO into Python on demand by CliMetNet
3. Training via Jupyter Hub on European Weather Cloud
4. Network search based on blueprints from other projects
5. Centralised solution to call neural network libraries within IFS
6. More code performance flexibility due to complementary work in the scalability project

Research highlights

Ackmann, J., P. D. Dueben, T. Palmer, & **P. Smolarkiewicz**, 2020: Machine-Learned Preconditioners for Linear Solvers in Geophysical Fluid Flows. https://arxiv.org/abs/2010.02866.

Adewoyin, R., P. D. Dueben, P. Watson, Y. He, and R. Dutta, 2020**:** Tru-net: A deep learning approach to high resolution prediction of rainfall. https://arxiv.org/abs/2008.09090.

Aires et al, 2021: Statistical approaches to assimilate ASCAT soil moisture information: Part I Methodologies and first assessment. *Q.J.R. Meteorol. Soc.* In review.

Baran, S, Baran, Á, Pappenberger, F, & **Ben Bouallègue, Z.**, 2020: Statistical post-processing of heat index ensemble forecasts: Is there a royal road? *Q.J.R. Meteorol. Soc.*, 146:3416-3434, https://doi.org/10.1002/qj.3853.

Barré, J. et al., 2020: Estimating lockdown induced European NO2 changes. *Atmos. Chem. Phys. Discuss.*, https://doi.org/10.5194/acp-2020-995.

Bonavita, M. & **P. Laloyaux**, 2020: Machine Learning for Model Error Inference and Correction. *Earth and Space Science Open Archive*, 10.1002/essoar.10503695.1.

Coughlan et al., 2021: Using machine learning to predict fire ignition occurrences from lightning forecasts. In press in *Met App*.

Geer, A. J., 2021: Learning earth system models from observations: Machine learning or data assimilation? In press in *Phil. Trans. A*.

Grönquist, Peter, et al., 2020: Deep Learning for Post-Processing Ensemble Weather Forecasts. In press in *Phil. Trans. A*.

Hewson, T. & **F. Pillosu**, 2020: A new low-cost technique improves weather forecasts across the world. arXiv:2003.14397v1.

Rasp, S., P. D. Dueben, S. Scher, J. A. Weyn, S. Mouatadid & **N. Thuerey**, 2020: WeatherBench: A benchmark dataset for data-driven weather forecasting. *Journal of Advances in Modeling Earth Systems*, **12**, e2020MS002203, https://doi.org/10.1029/2020MS002203.

Sherman Lo, Peter Watson, Peter Dueben, and Ritabrata Dutta, 2020: High-resolution probabilistic precipitation prediction for use in climate simulations, https://arxiv.org/abs/2012.09821.

A Roadmap for Machine Learning Activities at ECMWF

We are **hiring** **a machine learning scientist** as part of the AI4Copernicus project (closing 31st January)

The WMO will publish a **challenge for S2S predictions** in April 2021

Special issue on *Benchmark datasets and machine learning algorithms for Earth system science data* in ESSD (Martin Schultz, Amber Leeson and David Carlson) and GMD (Peter Dueben)

There will be many opportunities to engage with ECMWF, for example via **workshops on machine learning** with the Member and Co-operating states

"Following the steps outlined in this roadmap will enable ECMWF to prepare for evolving needs of scientists and analysts towards a more data-driven workflow and to support the Member and Co-operating States to make the most of new capabilities of machine learning as soon as possible."

Many thanks to the co-authors: Umberto Modigliani, Alan Geer, Stephan Siemen, Florian Pappenberger, Peter Bauer, Andy Brown, Martin Palkovič, Baudouin Raoult, Nils Wedi, Vasileios Baousis

Please let me know if you have any **questions:** peter.dueben@ecmwf.int // @PDueben

ECMWF | EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS

The strength of a common goal