npj | climate and atmospheric science

www.nature.com/npjclimatsci

MEETING REPORT OPEN

Check for updates

2022 ECMWF-ESA workshop report: current status, progress and opportunities in machine learning for Earth System observation and prediction

Massimo Bonavita[1], Rochelle Schneider[1,2✉], Rossella Arcucci[3], Matthew Chantry[1], Marcin Chrust[1], Alan Geer[1], Bertrand Le Saux[2] and Claudia Vitolo[2]

This report provides a summary of the main outcomes of the 3rd edition of the workshop on Machine Learning (ML) for Earth System Observation and Prediction (ESOP/ML4ESOP) co-organised by the European Centre for Medium-Range Weather Forecasts (ECMWF) and European Space Agency (ESA). The 4-day workshop was held on 14-17 November 2022 in hybrid format, with an in-person component at the ECMWF Reading site and an interactive online component, attracting a record number of submissions and over 700 registrations. The workshop aimed to document the current state-of-the-art, progress and challenges in the rapidly evolving field of the integration of ML technologies in ESOP and to provide a venue for discussion and collaboration for ESOP and ML specialists. The workshop was structured along five main thematic areas covering the principal components of standard ESOP workflows. Highlights from the presentations and a discussion of the most promising development directions from the workshop Working Groups in all the different thematic areas are provided in this Report.

npj Climate and Atmospheric Science (2023)6:87 ; https://doi.org/10.1038/s41612-023-00387-2

INTRODUCTION

The third edition of the ECMWF–ESA Workshop on Machine Learning for Earth Observation and Prediction (ML4ESOP) took place from 14 to 17 November 2022 at ECMWF Reading, UK (https://events.ecmwf.int/event/304/). After two exclusively online editions[1,2], this edition of the workshop was run in a hybrid format, hosting about 120 people on-site and large and active virtual participation (about 700 registered participants). The attendance numbers, together with a record number of 121 abstract submissions, confirm the large interest in machine learning (ML) for Earth system science applications and also the growing popularity of the ECMWF–ESA workshop series as a reference meeting and discussion venue in this area.

As it has become traditional in this series of workshops, two leading experts were invited to provide broad overviews of the state-of-the-art, current opportunities, and challenges in their field of expertise.

Prof. Stephen Penny's talk concentrated on the interactions and synergies between data assimilation (DA) and ML[3]. One fascinating suggestion was that current DA could leverage ML tools and ideas to greatly increase efficiency, which in turn could make more advanced algorithms affordable, e.g., allowing much larger ensemble sizes in ensemble DA and thus opening the way to more fully nonlinear/non-Gaussian assimilation methods.

Prof. Damien Borth's presentation focussed on recent advances in ML tools used in Earth Observation (EO) and Remote Sensing[4,5]. The main focus of the talk was on efficient representation learning, i.e., the increasingly sophisticated set of techniques that allows a ML model to automatically discover the representations needed for feature detection or classification from raw data. This is a field of research in ML which is particularly relevant for EO, where there are huge quantities of unlabelled, remotely sensed data which could potentially be used.

From both these talks, it was apparent that increasingly sophisticated ML techniques have further spread into research and operational practice in the Earth sciences and, more importantly, they are being tailored to this specific domain with compelling results.

The workshop was structured according to separate thematic areas (TA) designed to cover the main application areas of ML in EO, numerical weather prediction (NWP), and climate prediction:

- TA1: ML for Earth observations
- TA2: Hybrid-ML in data assimilation
- TA3: ML for model emulation and model discovery
- TA4: ML for user-oriented Earth science applications
- TA5: ML at the network edge and high-performance computing

The following sections will describe in more detail the presentations and discussions in the working groups on each TA.

TA1: ML FOR EARTH OBSERVATIONS

Current ML applications

The working group was chaired by Rochelle Schneider (ESA) and Alan Geer (ECMWF). Membership of the working group indicated a very broad interest in ML4EO that spans university, national meteorological services, and the private sector, with attendees drawn from Europe, America, Asia, Africa, and Australia. Application areas were similarly broad, including renewable energy, hydrology, clouds, environmental health, pollution, wildfires, urbanisation, geodesy, and crop classification.

[1]European Centre for Medium-Range Weather Forecasts, Reading, UK. [2]European Space Agency, Frascati, Italy. [3]Imperial College London, London, UK.
✉email: rochelle.schneider@esa.int

Limitations, opportunities, and challenges

Much of the discussion reflected how far EO applications of ML have progressed in terms of serving real applications. To be beneficial to society, operational AI tools should ideally be reproducible, scalable, maintainable, transferable, and explainable. They need to integrate well with existing non-Artificial Intelligence (AI) tools rather than compete, i.e., aiming to fully replace state-of-the-art methods. Further, the 'train once' approach often needs to be replaced by continuous learning and retraining from fresh observations.

The perennial question of finding enough training data brought up the contrast between applications that can use non-expert labelling, achievable through citizen science or gamification (e.g., zooniverse.org), applications that require domain experts for labelling, and those where 'ground truth' observations are required[6]. Developments in the wider AI community could help, including the idea of 'one shot' learning (a human only needs to see one Zebra), meta-learning (learning how to learn) and the possibility of using non-domain foundation models in EO (could a foundation model transition from cats and dogs to crop type identification?)

Another widely recognised issue was the choice between using in-house or cloud-based computing resources[7]. The latter are attractive for the easy access to sophisticated TPUs and GPUs to speed up training, and the possibility of rapidly scaling up applications thereafter. But issues were raised over the expense relative to in-house hardware, the dependence on private backends into which an application can become locked-in, a perceived lack of support and doubts over data protection and security[8].

Future directions

Including AI onboard satellites could speed up event detection (the idea of 'smart satellites') and could also help with data compression, prioritisation, and cooperation between satellites[9,10]. However, it is important that the full raw observational data is downlinked to earth and archived, possibly with less timeliness, in order to support future learning and development. The group also discussed federated learning, particularly for applications which rely on sensitive data (e.g., health), where AI is trained on dispersed data that is kept private.

TA2: HYBRID DATA ASSIMILATION—ML APPROACHES

Current ML applications

Rossella Arcucci (Imperial College London) and Marcin Chrust (ECMWF) co-chaired this working group, which explored the utilisation of hybrid ML and DA approaches. A very broad group of members coming from academia, industry, NWP centres, and research centres took part in an active discussion about the potential use of ML with DA from a perspective of improving DA modelling. Hybrid approaches based on ML and DA have become increasingly popular in the DA field, with applications spanning from using neural networks to emulate model components within DA to entirely replacing the well-established DA algorithms with ML-based emulators or bespoke techniques. The latter include extended Elman networks[11] that estimate posterior covariances and recurrent neural networks that replace 4D-Var with joint model and solver learning[12]. A large body of literature has been also dedicated to performing DA in latent spaces of a ML algorithm, balancing accuracy and computational cost[13]. The link between probabilistic ML approaches and differential equations is highlighted when the frameworks of DA and ML are combined from a Bayesian perspective. This equivalency, which demonstrates the parallels between the two areas, is presented formally in refs. [14,15].

Limitations, opportunities, and challenges

Learning full models or replacing DA algorithms with ML approaches is still considered challenging in operational settings given the difficulty posed by the high-dimensional nature of the systems involved. The development of hybrid models that combine physics-based models with statistical models has been proposed as an attractive alternative. It has been shown that statistical models that are used to correct physics-based models can be made situation dependent and trained within the framework of 4D-Var[16]. It was agreed that developing a common ML-DA framework would offer the potential to leverage the strengths of both approaches. ML, in particular, may present an opportunity to surpass the linearity and Gaussianity constraints imposed in current DA schemes utilised in numerical weather centres, while also significantly reducing the cost of the analysis process. The latter aspect is of growing importance as the resolution of the models and the analysis is increased in future.

Future directions

The combination of ML with DA advances the state-of-the-art of ML modelling in various fields and applications. The development trends and future challenges of this fast-growing field include learning state-observation mapping in DA or developing ML surrogates of dynamical systems assisted by DA. The working group participants reached a broad consensus that conventional DA methods can be used to improve ML algorithms, especially in addressing issues associated with noisy, incomplete or biased data. These hybrid models provide strengths in interpretability and noise reduction. Significant space for further breakthrough advances still exists, especially in applying these approaches in operational contexts.

TA3: ML FOR MODEL EMULATION AND MODEL DISCOVERY

Current ML applications

The working group was chaired by Massimo Bonavita and Matthew Chantry (both ECMWF). Assessing the state of the field for model emulation the group saw a full spectrum of approaches, starting from learning to emulate a component of a weather or climate model up to learning to emulate realistic weather models in their entirety. On the latter approach, significant progress had been observed in the past year, with claims that some ML models have become competitive with operational state-of-the-art models for deterministic predictions[17–19]. Model discovery is currently less common, but successful work on e.g., predicting rogue waves gave a blueprint for a successful application[20].

Limitations, opportunities, and challenges

Training of ML models was viewed as being overly reliant on mean-squared-error (or similarly constructed) loss functions, which has drawbacks as models trained towards this metric make cautious and overly smooth predictions which can be limiting in real world applications. Detailed discussions were held on the possibility of using Generative Adversarial Networks (GANs), diffusion models[21], and others. Also, the general approach of training ML models using probabilistic instead of deterministic loss functions was indicated as a promising way forward.

Future directions

From the presentations and working group discussions it was apparent that using ML for model emulation and, more broadly, for general forecasting purposes is developing rapidly and the entry in the field of major commercial players will further accelerate progress in this area. Whether these efforts will pose a fundamental threat to the traditional Numerical Weather and

Climate Prediction workflows based on physics-based weather emulators remains to be seen.

TA4: ML FOR USER-ORIENTED EARTH SCIENCE APPLICATIONS

Current ML applications

The working group was chaired by Claudia Vitolo and Bertrand Le Saux (both from ESA) and gathered a large and varied group of experts, as ML is now pervading all sectors of Earth System sciences and industry[22]. In the weather and climate domain, for instance, Deep Learning was reported to be used for nowcasting of precipitation[23], the detection of extreme weather events[24], post-processing of predictions (e.g., downscaling,[25], as well as for the analysis of climate and weather processes at a longer time scale. In environmental applications, participants reported to have used ML to get actionable information from EO data in a variety of fields, including public health, agriculture, environmental protection (on land and at sea) just to mention a few. Several compelling industrial applications were also mentioned and included during-flight planning (aviation sector), road maintenance planning (transportation sector), energy demand and distribution planning and (re)-insurance. Participants also briefly touched upon new emerging fields in which ML may deliver promising applications, for security or policymaking, in the near future which build on federated learning, onboard processing, digital twin technologies and quantum ML.

Limitations, opportunities, and challenges

A great challenge of ML in user-oriented Earth science applications is the lack of trust in black-box models and the lack of a common language between developers and domain experts. However, many participants were confident that the development of explainable AI[26], hybrid modelling, use of large pre-trained models, as well as strong community building will help to bridge the gap and lower the barrier to ML adoption. Technical issues such as costly compute resources and complex software usage are also deemed as critical and potentially limiting, and initiatives to provide accessible cloud computing, open and reusable source code are seen as potential enablers. Lastly, generative models (for image, text and data) are scrutinised with circumspection as they appear to have great potential for new applications[27] but in the meantime raise problems of transparency and ethics[28].

Future directions

Several perspectives emerged from the discussions and presentations. The development of new techniques and frameworks is considered highly impactful on society (for low-cost speed-up of complex numerical models, for fast simulations) and climate (for improving early warnings and identifying new sustainable environmental solutions for challenges affecting many sectors, e.g., energy). According to several participants, transformer models and explainable ML are deemed extremely promising to overcome the current lack of trust in black-box models[22,29]. AI is expected to be used more and more in extreme weather event predictions and digital twins modelling coupling various Earth system processes. In the longer term, there is a demand for investigating green computing, operationalisation, transparent ML and process understanding in Earth system sciences.

TA5: ML AT THE EDGE AND HIGH-PERFORMANCE COMPUTING

Current ML applications

During the workshop, novel computing with transformative power was a ubiquitous topic of discussion. With the rapid advances in computing technology, ML-enhanced high-performance computing has become an increasingly important tool in Earth science research, as illustrated by Carlos Alberto Gómez Gonzalez from Barcelona Supercomputing Centre in his talk on deep learning for empirical downscaling[30] used to obtain e.g., estimates of nitrogen dioxide or precipitation fields at fine resolutions. At the other extremity of the computer power spectrum, on-board processing with machine learning at the edge has already proved to be useful for Earth observation, as it enables real-time processing of satellite data and instantaneous response to events such as floods[10]. It also helps reduce the cost of data transmission, as the data is processed on the satellite itself before being transmitted to the ground.

Limitations, opportunities, and challenges

Modular computing environments, that is, systems which integrate different types of compute resources, are seen as a way to offer a degree of flexibility and scalability to large-scale computing applications. Thus, program parts of complex simulations can be distributed over several modules so that various hardware properties can be optimally leveraged. This also allows for the addition and removal of components to meet changing demands and requirements or integrate future technologies such as quantum computing or neuromorphic modules. More powerful supercomputers built on this principle might be the way to enter the exascale era of computing. And such capacity is likely to be required to face the new challenges of this century: running numerical models with enough precision to allow forecast of weather events at a local scale or supporting the development of digital twins to enable monitoring, forecasting and assessing the impact of climate change, as mentioned by Jacqueline Lemoigne from NASA in her talk, showing examples such as IDEAS[31].

Future directions

Many exciting perspectives were drawn and discussed. Distributed computing in space might offer the possibility to optimise the collaboration of small sensing satellites and satellites with computing payloads and enable cognitive cloud computing in space (C3S). Lisa Woerner from DLR highlighted the potential of quantum techniques for global Earth observation to climate change impact reduction. She stressed in particular the promises of quantum computing and quantum ML for which potential benefits could be gained from further exploration of under-utilised areas of machine learning, such as reinforcement learning[32]. Bertrand Le Saux from ESA detailed the ongoing efforts to bring the power of Quantum Computing to Earth observation and presented two fields of research on the use of quantum ML for such classical data: quantum kernels and hybrid classical-quantum neural networks. They already yield the first proofs of concepts for image classification on gate-based quantum computers[33,34] or quantum annealers[35] and analysis of time-series[36]. These hybrid quantum-classical architectures form the foundation of the next generation of ML that can be run on modular HPC.

CONCLUSION

We believe the workshop continues to be a valuable environment for fostering knowledge exchange and facilitating breakthrough discoveries in the field of ML4ESOP. The range of applications of ML technologies in ESOP is truly remarkable and keeps on growing (for example, the new application areas bridging AI with Quantum Computing and the use of AI solutions close the data, e.g., onboard orbiting satellites). Themes like reproducibility and interpretability of ML outputs, together with scalability and maintainability of ML technologies still dominate the discussion towards the operational application of ML solutions. From the working groups outcomes, a clear consensus emerged of the need and opportunity to leverage ML strengths to try to fill gaps in our knowledge-based models and improve perceived weaknesses in current operational methods. On the other hand, a philosophical and practical divergence of approaches is beginning to appear in the application of ML technologies in ESOP. On one side, a gradual, incremental adoption of ML solutions in established workflows, aimed at improving results and reducing computational costs while striving to maintain a more or less complete understanding of the modelled system. This is the avenue typically chosen by domain scientists and practitioners in ESOP. On the other side, a growing corpus of research aims to show the potential of AI/ML to disrupt traditional practices through end-to-end, completely data-driven ML/AI solutions. In these applications, the modelled system is seen as a black (or grey) box, but the computational efficiencies are compelling and the quality of the predictions is becoming competitive with the state-of-the-art. This is the avenue favoured by some ML researchers approaching the ESOP world. It will be fascinating to see which approach will turn out to be more fruitful in the long run, but it can already be said that this dynamic has set in motion a far-reaching debate in the ESOP community on current methodologies and their long-term sustainability.

Received: 7 March 2023; Accepted: 16 May 2023;
Published online: 12 July 2023

REFERENCES

1. Bonavita, M. et al. Machine learning for earth system observation and prediction. *Bull. Am. Meteorol. Soc.* **102**, E710–E716 (2021).
2. Schneider, R. et al. ESA-ECMWF Report on recent progress and research directions in machine learning for Earth System observation and prediction. *npj Clim. Atmos. Sci.* **5**, 51 (2022).
3. Penny, S. G. et al. Integrating recurrent neural networks with data assimilation for scalable data-driven state estimation. *J. Adv. Model. Earth Syst.* **14** (2022). https://doi.org/10.1029/2021MS002843.
4. Scheibenreif, L., Hanna, J., Mommert, M. & Borth, D. Self-supervised vision transformer for land-cover segmentation and classification. *CVPR Earth Vision-Workshop* pp. 1422–1431 (2022).
5. Schürholt, K., Knyazev, B., Giró-i-Nieto, X. & Borth, D. Hyper-representations as generative models: sampling unseen neural network weights. Neural Information Processing Systems (NeurIPS) (2022).
6. Dueben, P. D. et al. Challenges and benchmark datasets for machine learning in the atmospheric sciences: definition, status, and outlook. *Artif. Intell. Earth Syst.* **1**, e210002 (2022).
7. Wagemann, J., Siemen, S., Seeger, B. & Bendix, J. A user perspective on future cloud-based services for Big Earth data. *Int. J. Digital Earth* **14**, 1758–1774 (2021). **12**.
8. Gomes, V. C. F., Queiroz, G. R. & Ferreira, K. R. An overview of platforms for big earth observation data management and analysis. *Remote Sens.* **12**, 1253 (2020).
9. Furano, G. et al. Towards the use of artificial intelligence on the edge in space systems: challenges and opportunities. *IEEE Aerosp. Electron. Syst. Mag.* **35**, 44–56 (2020).
10. Mateo-Garcia, G. et al. Towards global flood mapping onboard low cost satellites with machine learning. *Sci. Rep.* **11**, 7249 (2021).
11. Boudier, P., Fillion, A., Gratton, S., Gürol, S. & Zhang, S. DAN-An optimal Data Assimilation framework based on machine learning Recurrent Networks - arXiv preprint arXiv:2010.09694 (2020).
12. Fablet, R. et al. Learning variational data assimilation models and solvers. *J. Adv. Modeling Earth Syst.* **13**, e2021MS002572 (2021).
13. Arcucci, R. Data Learning for more reliable digital twins. 3rd ECMW-ESA Workshop on Machine Learning for Earth System Observation and Prediction (2022). Available at: https://events.ecmwf.int/event/304/contributions/3630/attachments/2153/3812/ECMWF-ESA-ML_Arcucci.pdf
14. Geer, A. J. Learning earth system models from observations: machine learning or data assimilation? *Philos. Trans. R. Soc. A* **379**, 20200089 (2021).
15. Bocquet, M., Brajard, J., Carrassi, A. & Bertino, L. Bayesian inference of chaotic dynamics by merging data assimilation, machine learning and expectation-maximization. *Found. Data Sci.* **2**, 55–80 (2020).
16. Farchi, A., Chrust, M., Bocquet, M., Laloyaux, P. & Bonavita, M. Online model error correction with neural networks in the incremental 4D-Var framework. arXiv preprint arXiv:2210.13817 (2022).
17. Pathak, J. et al. Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators. arXiv preprint arXiv:2202.11214 (2022).
18. Bi, K. et al. Pangu-weather: A 3d high-resolution model for fast and accurate global weather forecast. arXiv preprint arXiv:2211.02556 (2022).
19. Lam, R. et al. GraphCast: Learning skillful medium-range global weather forecasting. Arxiv preprint (2022). https://arxiv.org/abs/2212.12794
20. Häfner, D., Gemmrich, J. & Jochum, M. Real-world rogue wave probabilities. *Sci. Rep.* **11**, 10084 (2021).
21. Dhariwal, P. & A. Nichol. Diffusion Models Beat GANs on Image Synthesis. arXiv:2105.05233 (2021).
22. Singh, M. et al. Artificial intelligence and machine learning in earth system sciences with special reference to climate science and meteorology in South Asia. *Curr. Sci.* **122**, 1019–1030 (2022).
23. Chen, L., Cao, Y., Ma, L. & Zhang, J. A deep learning-based methodology for precipitation nowcasting with radar. *Earth Space Sci.* **7**, e2019EA000812 (2020).
24. Toğaçar, M., Ergen, B. & Cömert, Z. Detection of weather images by using spiking neural networks of deep learning models. *Neural Comput Applic* **33**, 6147–6159 (2021).
25. Harris, L., McRae, A. T. T., Chantry, M., Dueben, P. D. & Palmer, T. N. A generative deep learning approach to stochastic downscaling of precipitation forecasts. *J. Adv. Modeling Earth Syst.* **14**, e2022MS003120 (2022).
26. Roscher, R., Bohn, B., Duarte, M. F. & Garcke, J. Explainable Machine Learning for Scientific Insights and Discoveries. *IEEE Access* **8**, 42200–42216 (2020).
27. Castillo-Navarro, J., Le Saux, B., Boulch, A. & Lefèvre, S. Energy-Based Models in Earth Observation: From Generation to Semisupervised Learning. *IEEE Trans. Geosci. Remote Sens.* **60**, 1–11 (2022).
28. Kochupillai, M., Kahl, M., Schmitt, M., Taubenböck, H. & Zhu, X. X. Earth Observation and Artificial Intelligence: Understanding emerging ethical issues and opportunities. *IEEE Geoscience and Remote Sensing Magazine* (2022). https://doi.org/10.1109/MGRS.2022.3208357.
29. Amanambu, A. C., Mossa, J. & Chen, Y.-H. Hydrological Drought Forecasting Using a Deep Transformer Model. *Water* **14**, 3611 (2022).
30. Gomez Gonzalez, C. A. DL4DS-Deep Learning for empirical DownScaling. arXiv e-prints, (2022) https://arxiv.org/abs/2205.08967
31. Huang, T., et al. Integrated Digital Earth Analysis System (IDEAS). AGU Fall Meeting Abstracts (2021)
32. Saggio, V. et al. Experimental quantum speed-up in reinforcement learning agents. *Nature* **591**, 229–233 (2021).
33. Sebastianelli, A., Zaidenberg, D. A., Spiller, D., Le Saux, B. & Ullo, S. L. On Circuit-based Hybrid Quantum Neural Networks for Remote Sensing Imagery Classification, IEEE J. Sel. Topics in Earth Obs. and Rem. Sens., **15**, pp. 565–580 (2021). https://doi.org/10.48550/arXiv.2109.09484.
34. Chang, S.-Y., Vallecorsa, S., Grossi, M. & Le Saux, B. Hybrid Quantum-Classical Networks for Reconstruction and Classification of Earth Observation Images. 21st Int. Workshop on Advanced Computing and Analysis Techniques (ACAT), *.Physics Research* pp. 4907–4910 (2022).
35. Delilbasic, A., Le Saux, B., Riedel, M., Michielsen, K. & Cavallaro, G. A Single-Step Multiclass SVM Based on Quantum Annealing for Remote Sensing Data Classification. IEEE J. Sel. Topics in Earth Obs. and Rem. Sens., in press (2023).
36. Siemaszko, M., McDermott, T., Buraczewski, A., Le Saux, B. & Stobińska, M. Rapid Training of Quantum Recurrent Neural Networks, Quantum Technologies in Machine Learning (2022). https://doi.org/10.48550/arXiv.2207.00378.

ACKNOWLEDGEMENTS

The organising committee wishes to thank all the speakers and participants at the 3rd workshop edition. The session chairs would also like to thank the participants who attended the hybrid working groups. We are also grateful for the help and support received from the communication teams from both organisations.

AUTHOR CONTRIBUTIONS

M.B. and R.S. defined the questions that structured the basis of the manuscript. M.B. took the lead in writing the manuscript. All authors contributed to writing and providing critical feedback to the manuscript.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Correspondence and requests for materials should be addressed to Rochelle Schneider.

Reprints and permission information is available at http://www.nature.com/reprints

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.