Second white paper on community guidelines on the use, value and applicability of emerging technologies in climate and weather applications

Deliverable D2.7

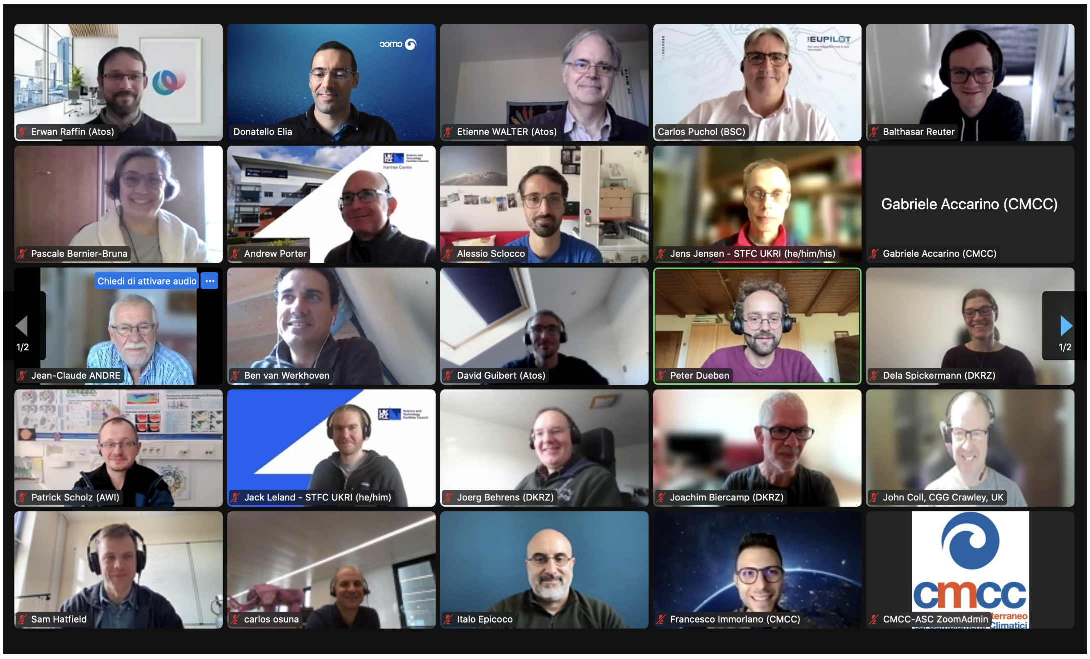

The project Centre of Excellence in Simulation of Weather and Climate in Europe Phase 2 (ESiWACE2) has received funding from the European Union's Horizon 2020 Research and Innovation Programme under Grant Agreement No 823988.

About this document:
Work package in charge: WP2 - Establish, evaluate and watch new technologies for the community
Actual delivery date for this deliverable: 21st November 2022
Dissemination level: PU (for public use)

Lead authors:
Giovanni Aloisio (CMCC)
Donatello Elia (CMCC)

Other contributing authors:
Gabriele Accarino (CMCC)
Graham Riley (University of Manchester)
Mike Ashworth (University of Manchester)
Workshop speakers (for the abstracts): see the list of speakers at the end of the document

Project internal reviewer:
Peter Dueben (ECMWF)

Contact details:
Project Office: esiwace@dkrz.de
Visit us on: https://www.esiwace.eu

Access our documents in Zenodo:
https://zenodo.org/communities/esiwace

Contents

1 Abstract /publishable summary

This second white paper provides an overview of trending and emerging technologies related to climate and weather applications, focusing on novel extreme-scale systems, programming models for heterogeneous computing and machine learning (ML). It updates the findings highlighted in the first white paper by providing a summary of relevant efforts available in literature, as well as some early results from these fields.

Tracking of state-of-the-art solutions and technologies has been carried out during the ESiWACE2 project to identify recent and relevant work focusing on trends in High Performance Computing (HPC), use of hardware accelerators and machine learning for weather and climate modelling. The document summarises the main results from this tracking activity providing pointers and references to interesting articles.

In order to track some of the earliest work in the field, a second *workshop on Emerging Technologies for Weather and Climate Modelling* was held on the 7th of October as a virtual event. The workshop organised in the frame of the ESiWACE2 project aimed to bring together HPC experts from academia and industry and climate/weather scientists from the EU and USA to discuss about some of the latest developments, research opportunities and efforts. The workshop agenda was organised around three topical sessions, for a total of 14 talks, about: *European Exascale hardware* (Session 1), *Programming models and hardware interplay* (Session 2) and *Machine Learning* (Session 3). More than 130 people from 22 countries world-wide registered to the workshop with Germany, UK, Italy, France and USA being the most represented countries.

2 Conclusions & Results

The overarching goal of this work is to raise awareness of the community and share experiences on the use of emerging technologies and solutions in the field of climate and weather research. The document provides an overview of some of the current and trending efforts in HPC, programming models and ML for weather and climate applications.

In particular, the workshop on emerging technologies reported in this document brought together HPC experts, supercomputing centres representatives and climate scientists to discuss about the latest developments in the aforementioned fields. EU industry partners from the ESiWACE2 project (i.e., ATOS) have been involved in the organisation of the workshop to further foster the discussion between technology developers and the research community.

As a result of this activity, some key trends have been identified:

- The infrastructure of the newest supercomputers being built is designed to support more data-driven workloads and encompass heterogeneous architectures making use of a large amount of accelerated hardware (e.g., GPUs and FPGAs);
- As the machines grow in performance capabilities and size, energy efficiency has become a central factor in the design of these systems;
- High-level programming models, abstractions and Domain Specific Languages (DSLs) are emerging as solutions to better support performance portability of climate/weather models on top of heterogeneous infrastructures;
- Machine learning approaches are now part of the HPC workload and are currently explored to address different aspects of weather and climate modelling such as emulation of numerical components, downscaling, prediction of extreme events, etc.

3 Project objectives

This deliverable contributes directly and indirectly to the achievement of all the macro-objectives and specific goals indicated in section 1.1 of the Description of the Action:

Macro-objectives	**Contribution of this deliverable**
(1) Enable leading European weather and climate models to leverage the available performance of pre-exascale systems with regard to both compute and data capacity in 2021.	X
(2) Prepare the weather and climate community to be able to make use of exascale systems when they become available.	X

Specific goals in the workplan	**Contribution of this deliverable**
Boost European climate and weather models to operate in world-leading quality on existing supercomputing and future pre-exascale platforms	
Establish new technologies for weather and climate modelling	X
Enhance HPC capacity of the weather and climate community	
Improve the toolchain to manage data from climate and weather simulations at scale	
Strengthen the interaction with the European HPC ecosystem	X
Foster co-design between model developers, HPC manufacturers and HPC centres	X

4 Detailed report on the deliverable

In the last few years, in the wake of exascale era, the HPC hardware ecosystem has radically changed encompassing more heterogeneous computing architectures jointly with technologies for data-driven applications. These new infrastructures open the field to unprecedented opportunities for scientific computing, besides introducing several challenges due to their increased complexity (Reed and Dongarra, 2015).

In the context of the ESiWACE CoE, a tracking activity has been carried out to identify some of the most recent developments and open challenges related to emerging technologies with a focus on those more relevant for the climate modelling and weather forecasting communities. In particular, in order to gain insights into the most cutting-edge activities, besides tracking the advancements available from literature, the task also organised a series of community-oriented workshops.

The rest of the document first presents an overview of the key state-of-the-art advancements related to HPC infrastructures, heterogeneous architectures and machine learning (section 4.1), followed by the description of the second workshop on emerging technologies organised in October 2022 (in section 4.2).

4.1 Technology tracking for climate modelling and weather forecast

4.1.1 Technology trends in High Performance Computing

The compute performance available to a wide range of scientific applications, including weather forecasting and climate prediction, through the widespread use of High Performance Computing (HPC) has increased markedly for many decades. This increase has followed an exponential trend, doubling approximately every eighteen months (Reed and Dongarra, 2015). Originally this was described by Moore's Law, which gave a prediction for the exponential increase in the number of gates on a chip (Moore, 2006). Yet, as the increase in HPC system performance now owes more to increased exploitation of parallelism than to increased gate density, it should be regarded as only a partial consequence of Moore's Law. Therefore, the end of Moore's Law (Markov, 2014) signals the end of only one opportunity for increased HPC performance.

The potential benefits of increased HPC performance in the weather and climate domain have been widely discussed (Neumann et al., 2019), giving opportunities for greater forecast accuracy and improvements in the quality of climate prediction through increased grid resolution, improvements in physics parameterisations, replacement of parameterisations by accurate modelling, greater use of ensemble forecasting and statistical techniques, etc. Delivering such benefits requires pushing HPC system performance, currently available at the 10s-100s Pflop/s, to the exascale (10^{18} flop/s) regime[1]. In addition to the drive for performance, the desire to limit costs and reduce climate impacts also mean that new performance improvements must be met with a reduction, or at least no increase, in power consumption, leading to a new focus on power efficiency[2], expressed as Gflop/s/W.

Performance and power efficiency goals will therefore only be met through the use of new and possibly disruptive technologies, rather than scaling up existing hardware. This process started more than a decade ago with the appearance of the first supercomputers with heterogeneous architectures (i.e, CPU+GPU) in the *TOP500* list (Kindratenko and Trancoso, 2011). The trend in the adoption of accelerated architectures in HPC has continued to increase ever since and, in the latest TOP500 list from June 2022[3], most of the top supercomputers integrate some type of hardware accelerator (mainly GPUs) to push the achievable performance.

[1] https://www.top500.org/
[2] https://www.top500.org/lists/green500/
[3] https://top500.org/lists/top500/2022/06/

In this *"race to exascale"*, investments are being made by different nations worldwide showing the strategic relevance of HPC for advancing research (Dongarra et al., 2019). Europe is also investing in building its own infrastructure of extreme-scale supercomputers under the *EuroHPC Joint Undertaking* (EuroHPC JU)[4]. The three largest systems currently installed (or under completion) procured through this initiative are *LUMI* at CSS[5], *Leonardo* at CINECA[6] and *MareNostrum 5* at BSC[7], which have been presented during the workshop (section 4.2.2). These are all pre-exascale systems featuring large numbers of GPU, besides CPU, that will allow supporting climate research applications, among others. In particular LUMI, the machine that was completed first, is currently third in the TOP500 list with a High Performance Linpack score of *151.9 Pflop/s*.

The changes in these new supercomputers will have an impact on scientific applications programming, and application developers will have to manage the transition from traditional to heterogeneous hardware platforms, which may have radically different and more complex architectures. Software portability, achieving correct results across different platforms, is one aspect of this, but performance benefits will only be delivered if applications achieve the target performance on the new hardware – performance portability (Deakin et al., 2019).

Current programming models have been successful for various compute-intensive applications in delivering performance on systems with millions of threads of execution to achieve Pflop/s of sustained performance. However, with the advent of exascale systems, efficient exploitation of the infrastructure will become more challenging. The leading system from the last TOP500 list is the *Frontier* machine installed at the Oak Ridge National Laboratory (ORNL) in the US, the first system to break the exascale wall with *1.102 Exaflop/s* on High Performance Linpack using 8,730,112 cores including 37,632 GPUs[8] [9]. A key question therefore is whether the same programming models will provide performance portability to systems of the future or emerging Domain Specific Languages (DSLs), such as those currently explored in the ESiWACE2 project, will fulfil their promise of separating the expression of algorithms from the (heterogeneous) hardware on which they are executed. Alternatively, will radical hardware changes necessitate the introduction of radically different programming models and what steps will be required to port existing applications to those new models?

One such potentially disruptive technology for HPC is the use of reconfigurable hardware in which the very circuits that define the low-level hardware architecture, which are burned into the silicon during manufacture for conventional CPUs and GPUs, are able to be reconfigured by software. The leading type of reconfigurable chip is the Field Programmable Gate Array (FPGA) in which the chip is provided with large numbers of simple logic elements of different types and a bitmap (known as a bitstream) is used to define how the elements are connected to construct the desired architecture. This allows far greater flexibility enabling the programmer to tailor the hardware architecture to the application. But flexibility comes at a cost to the application developer. Programming environments for reconfigurable computing expose a huge multi-dimensional parameter space for optimisation, making it difficult to home in on the best set of design choices for a particular application running on a particular architecture.

The EU funded project *EuroExa* project[10], full title *Co-designed Innovation and System for Resilient Exascale Computing in Europe: From Applications to Silicon*, proposes an HPC architecture which is both scalable to exascale performance levels and delivers world-leading power efficiency. This is achieved through the use of

[4] `https://eurohpc-ju.europa.eu/`

[5] `https://www.lumi-supercomputer.eu/lumi_supercomputer/`

[6] `https://leonardo-supercomputer.cineca.eu/`

[7] `https://eurohpc-ju.europa.eu/marenostrum5-new-eurohpc-world-class-supercomputer-spain-2022-06-16_en`

[8] `https://top500.org/news/ornls-frontier-first-to-break-the-exaflop-ceiling/`

[9] `https://spectrum.ieee.org/frontier-exascale-supercomputer`

[10] `https://www.euroexa.eu/` EuroExa has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement no 754337.

low-power ARM processors together with closely coupled FPGA programmable components. EuroExa combines state-of-the-art computing components using a groundbreaking system architecture, which applies the design flexibility of *UNIMEM* architecture from the EuroServer project (Durand et al., 2014), delivers high levels of performance to the selected applications, and balances compute resources with the resource demands of applications. Through co-design between the enabling technologies, the system software and the applications, EuroExa is delivering an innovative solution that achieves both extreme data processing and extreme computing. This solution will be demonstrated through the design, construction, testing and evaluation of three testbed systems throughout the duration of the project. This will enable EuroExa to deliver a recipe for the creation of an exascale computer by the end of the project.

4.1.2 Exploration of accelerators use in weather and climate

The increasing complexity of the supercomputing architectures, as described in the previous session, calls for higher-level programming models able to hide the complexity of the infrastructure to scientific code developers and enable performance portability across heterogeneous computing hardware. *MPI* and *OpenMP*, which remain among the most used programming models for parallel coding, have introduced in the last years some capabilities to enable interactions with GPUs: e.g., offloading parallel code on GPU devices[11] or through GPU-aware MPI libraries implementations (Khorassani et al., 2019). Besides these traditional programming models, new ones emerged in the last decade for supporting scientific computing on heterogeneous hardware (mainly GPUs), with different approaches followed by the US and EU communities. Some examples of these include the *Raja*[12] (Beckingsale et al., 2019) and *Kokkos* (Edwards et al., 2014) programming model abastractions from the US community and the *PSyclone*[13] and *DAWN*[14] climate/weather DSLs explored in the ESiWACE2 project. Other similar efforts for climate and weather modelling include the *GridTools for Python* (GT4Py)[15], i.e., a Python-based DSL, and the *CLAW* DSL for Fortran code (Clement et al., 2018).

These new programming abstractions can allow easier porting of climate and weather models on different processing devices, by providing higher-level and hardware-independent languages promoting a separation of concern between science code and parallel computing (Adams et al., 2019; Bertagna et al., 2019; Ben-Nun et al., 2022). Some results of using these high-level programming models have also been presented during the Emerging Technology workshop (in section 4.2.3).

Another relevant area of research is related to the use of FPGAs for scientific computing. The HPC community's interest in FPGAs as accelerators has been renewed due to the introduction of the *High-Level synthesis* tools (HLS). HLS tools hide the complexity of FPGA programming through raising the abstraction level for programmers. They offer environments where traditional HPC programmers can use high-level languages such as C/C++ and *OpenCL* to implement HPC application kernels on FPGAs. For the traditional HPC scientific programmer, the appropriate HLS environment choice implies trade-offs between the achievable performance and programmer effort. This is related to the level of detailed knowledge of FPGA hardware and performance improvement techniques assumed by the HLS environments.

A paper submitted to the *27th International Conference on Parallel & Distributed Processing Techniques & Applications* (PDPTA'21) presented a comparative study between three HLS programming methodologies: *Xilinx Vivado HLS* and *Xilinx SDSoC* using both OpenCL and C++, all targeting the X*ilinx Zynq UltraScale+ MPSoC ZCU102* (Alghamdi et al., 2021). Based on an existing low-level Vivado HLS design of a

[11] `https://www.openmp.org/updates/openmp-accelerator-support-gpus/`
[12] RAJA Performance Portability Layer. `https://github.com/LLNL/RAJA`
[13] `https://psyclone.readthedocs.io/en/stable`
[14] `https://github.com/MeteoSwiss-APN/dawn`
[15] `https://github.com/GridTools/gt4py`

finite-element kernel in the context of a realistic *LFRic* Climate model mini-app, the paper made a comparison between the programming techniques, effort and resulting performance of the implementations developed using the higher level of abstraction provided by SDSoC C/C++ and OpenCL. In addition, a comparative analysis was provided to illustrate the insights that led to the different design choices, scaling behaviour and peak performances obtained. It was found that Vivado HLS provides the highest performance due to the programmer's ability to exploit low-level FPGA features in the manual construction of the hardware system design, but near equivalent solutions can be obtained with OpenCL and C++ with a higher-level view of the FPGA hardware, including automatic generations of the system's design, and, hence, less programmer effort. This work was a collaboration at UNIMAN between an externally-funded PhD student, Moteb Alghamdi, and Graham Riley and Mike Ashworth in the ESiWACE2 project (Alghamdi et al., 2021).

In addition, a bibliography of papers (joint work with IS-ENES3) about the use of FPGAs in weather and climate modelling has been produced[16].

Within the EuroExa project, and in order to demonstrate the efficacy of the design, the EuroExa partners have assessed performance using a rich set of applications. One such application is the new weather and climate model, LFRic (named in honour of Lewis Fry Richardson), which is being developed by the Met Office and its partners for operational deployment in the middle of the next decade (Adams et al., 2019). High quality forecasting of weather and climate on global, regional and local scales is of great importance to a wide range of human activities, and the exploitation of latest developments in HPC has always been of critical importance to the weather forecasting and climate research communities.

The first steps in porting the LFRic Weather and Climate model to the FPGAs of the EuroExa architecture are described by (Ashworth et al., 2019). This paper describes the use of Vivado High Level Synthesis to implement a matrix-vector kernel from the LFRic code on a Xilinx UltraScale+ development board containing an *XCZU9EG* Multi-Processor System-on-a-Chip. It further details porting of the code, discusses the optimization decisions and reports performance of 5.34 Gflop/s with double precision and 5.58 Gflop/s with single precision.

(Ashworth et al., 2019) go on to discuss sources of inefficiencies, comparisons with peak performance, comparisons with CPU and GPU performance (w.r.t. power and price), comparisons with published techniques, and with published performance, and conclude with some comments on the prospects for future progress with FPGA acceleration of the weather forecast model.

The realisation of practical exascale-class high-performance computing systems requires significant improvements in the energy efficiency of such systems and their components. This has generated interest in computer architectures which utilise accelerators alongside traditional CPUs. FPGAs offer huge potential as an accelerator which can deliver performance for scientific applications at high levels of energy efficiency. The EuroExa project has been developing and building a high-performance architecture based upon ARM CPUs with FPGA acceleration targeting exascale-class performance within a realistic power budget.

4.1.3 Bibliography study of ML models for climate and weather

ML use is becoming widespread in several scientific domains, such as for climate and weather, since it can provide opportunities to improve modelling accuracy and efficiency. As it can be observed from the analysis of the literature, including the workshop talks, several research efforts are being carried out by the community. This section provides an overview of some of the state-of-the-art works in this field.

In particular, a survey of machine learning technologies applied to weather and climate modelling over the past ten years or so was undertaken in 2019 in conjunction with an MSc student from the University of Manchester, Yu Qi and a contribution from IS-ENES3 (Qi, 2019). The survey found over one hundred papers

[16] `https://drive.google.com/file/d/13Fw5M663whlKF4LBzm8TON-jEnxIrTzO/view?usp=sharing`

and resulted in a searchable spreadsheet-based database which classified the papers in terms of scientific application focus, ML technology used, the institution undertaking the research, etc. The study also included the analysis of the geographical distribution of the institutions. The database is available as a Google spreadsheet at: ML in Weather and Climate Database, `https://docs.google.com/spreadsheets/d/1vQ0CiSxlBx-_sfg6bMuXDS2ZoZkEBJaQ/edit#gid=526021773`[17].

An additional analysis of the literature was carried out to highlight some of the main applications of ML models in the domain. The analysis focused on three key usage scenarios: downscaling, extreme events studies, physics-guided ML models. The following subsections discuss some of these efforts.

ML-based Downscaling approaches

Downscaling is a procedure that allows making predictions at local scales, starting from climatic field information available at large scale. Recent advances in deep learning (*DL*) provide new insights and modelling solutions to tackle downscaling-related tasks by automatically learning the coarse-to-fine grained resolution mapping (Accarino et al., 2021).

The climatic fields at large scales are resolved by the Global Climate Models (*GCMs*), whereas the fields at local scales are resolved by Regional Climate Models (*RCMs*). Mapping this information is crucial to understand the local climate dynamics, which are often critical for assessing the impacts of a changing climate on society (Baño-Medina et al., 2018; Vandal et al., 2019). Unlike traditional dynamical approaches, statistical ones involve learning the empirical statistical relationships between coarse GCM outputs and High-Resolution (*HR*) products, including in-situ observations (Baño-Medina et al., 2018; Vandal et al., 2019; Baño Medina et al., 2020; Sachindra et al., 2018).

Among statistical approaches, recent advances in ML and DL provide new insights and modelling solutions to tackle downscaling-related tasks by automatically learning the coarse-to-fine resolution mapping. These algorithms may be trained on a historical dataset in order to learn the mapping between the low-resolution map of a given variable, such as temperature, precipitation, etc., and its high-resolution counterpart targeting the selected geographical domain (e.g., the European domain). These variables are typically treated as images, and auxiliary physical information could be further used as input by stacking the corresponding maps on the original image as additional channels. This could help the algorithm to better learn the underlying mapping between coarse and fine resolutions.

The dataset may be built through gathering and pre-processing analysis and reanalysis model products, such as *ERA-Interim*[18] or *ERA5* (Hersbach et al., 2020), but it can also be extended to satellite imagery and retrievals (Alemohammad et al., 2018).

LASSO regression was used in (Gao et al., 2014) for downscaling precipitation, whereas the random forest (*RF*) algorithm was used in (Bartkowiak et al., 2019) for land surface temperature. Artificial Neural Networks (*ANN*s) have been used in (Salimi et al., 2019) to perform precipitation downscaling. Furthermore, a Back-Propagation Neural Network (*BPNN*) and Support Vector Machine (*SVM*) fusion approach was adopted in (Min et al., 2020) to downscale precipitation. Several works moved towards deep architectures in the context of DL, especially concerning Long-Short Term Memory (*LSTM*) networks, Convolutional Neural Networks (*CNN*s) and Generative Adversarial Networks (*GAN*s). Since their introduction presented in (Hochreiter and Schmidhuber, 1997), LSTMs have been proven to be suitable for recovering and bridging information arbitrarily far in time, while avoiding the vanishing gradient problem. LSTMs have also been widely adopted for time-series related problems and, in the context of climate downscaling, for statistical downscaling of

[17] open with Google sheets to get access to the search features in the top row of the spreadsheet
[18] `https://apps.ecmwf.int/datasets/data/interim-full-daily/levtype=sfc/`

precipitation (Misra et al., 2018) and rainfall forecasting (Tran Anh et al., 2019) in combination with Feed Forward Neural Networks (FFNNs). CNNs, due to their ability to deal with spatio-temporal multi-dimensional structures, have been demonstrated to be particularly suitable for accomplishing Super-Resolution (*SR*) tasks. Several attempts to use deeper architectures have been proposed (Dong et al., 2014, 2016; Kim et al., 2016; Johnson et al., 2016; Kim et al., 2018; Park et al., 2019) for the extraction of high-level image characteristics and the downscaling of climatic fields. A deep neural network based on a CNN and a LSTM recurrent module was proposed in (Miao et al., 2019) to estimate precipitation based on well-resolved atmospheric dynamical fields. A novel architecture, named *DeepSD*, based on the super resolution framework, was presented in (Vandal et al., 2017, 2018) for downscaling precipitation fields, and a CNN-based approach was proposed in (Pan et al., 2019) as an alternative solution to the existing precipitation-related parameterisation schemes for the numerical precipitation estimation.

A CNN model for downscaling the occurrence of precipitation was also proposed in (Baño-Medina et al., 2018), whereas different configurations of CNN were adopted in (Sun and Lan, 2021) to downscale daily temperature and precipitation over China. A competitive DL framework based on a CNN was presented in (Huang, 2020) for downscaling temperature and precipitation, and it performed particularly well in generating spatio-temporal details at very fine-grid scales. A U-Net-type CNN was also used in (Kern et al., 2020) to learn a one-to-one mapping of low-resolution (input) to high-resolution (output) wind fields simulation data, and a conditional variational autoencoder (based on CNN) was exploited for learning the conditional distributions from data, assessing multimodalities and uncertainties. A CNN was adopted in (Shi, 2020) to perform smart dynamical downscaling (*SDD*) for extreme precipitation events, whereas a surrogate model, based on a Deep CNN (*DCNN*), was evaluated in (Sekiyama, 2020) for surface temperature, and was found to estimate image details that were not retained in the inputs. Recently, remarkable results were reported in several studies exploiting GANs for SR tasks in climate science. An Enhanced Super-Resolution GAN (*ESRGAN*) (Wang et al., 2019) was adopted and presented in (White et al., 2019) to downscale wind speeds from a coarse grid, capturing high frequency power spectra and high order statistics in the dataset, thus generating images of superior visual quality compared to the SR-CNN. A novel method (ClimAlign) was introduced in (Groenke et al., 2021) for unsupervised, generative downscaling of temperature and precipitation based on normalising flows for variational inference. In other works, a GAN was used in (Leinonen et al., 2021) for downscaling Time-Evolving Atmospheric Fields, while in (Harris et al., 2022) a GAN was used for stochastic downscaling of precipitation forecasts. Further works (Baño Medina et al., 2020; Mendes and Marengo, 2010; Mouatadid et al., 2017; Chang et al., 2018; Liu et al., 2018; Sharifi et al., 2019; Höhlein et al., 2020; Xu et al., 2020; Li et al., 2020) opted for downscaling based on ML and DL, and they helped assess both strengths and weaknesses of such methods. The results reported in these studies show that downscaling models based on ML allow better performance with respect to the other statistical approaches presented in (Sachindra et al., 2013; Goly et al., 2014; Duhan and Pandey, 2015).

The main challenges of the statistical downscaling methods, which ML and DL belong to, are:

- generalising well to data in which there is an emergent climate change signal, that might have not been seen during the training process on historical data
- generalising also in other spatial domains not seen during the training process (Relocatable Downscaling)

Additionally, the pre-trained ML and DL models on historical data could be used to infer long-term climate patterns on regional and local scales by using projection data as input. Different Shared Socioeconomic Pathways (*SSP*s) may be used to account for different levels of radiative forcing emissions according to different changes in global society, demography and economy.

ML for Extreme Weather Events

In climate science, occurrences of a significant change in the atmospheric and/or oceanographic conditions in a specific time instant is called an "event" (KanimozhiSelvi and Sowmiya, 2019). In case of a very small probability of occurrence, scientists talk about Extreme Weather Events (*EWE*s). They cause catastrophic weather phenomena - such as short-term heavy rainfall, thunderstorms, gales, tornadoes, hail, etc. - that are usually sudden, have a short duration and manifest at local scales (Fang et al., 2021). Between 2000 and 2020, the cumulative number of disasters worldwide reached 13,345, and more than 1.5 million people died (Yan, 2020): this is why a lot of research has been devoted to EWEs in the scientific community.

Tropical Cyclones (*TC*s), also popularly known as hurricanes or typhoons, are among the most spectacular and deadly geophysical phenomena (Emanuel, 2003). Generally, the uncertainty of weather and the complexity of atmospheric processes make extreme weather an important and difficult meteorological problem (Fang et al., 2021). Deep Learning models are able to automatically extract and learn intricate patterns from large datasets and to exploit this information to make subsequent predictions. Indeed, they have already proved to be efficient tools for tackling several climate science-related tasks, like TCs detection.
Detection starts with the classification of the presence of a cyclonic phenomenon in a particular time instant and on gridded climatic fields that result significant to its cyclogenesis. If the cyclone is classified as present, the following step is the identification of the geographical coordinates of its center in terms of latitude and longitude. Since the TC could potentially occur in different positions of the considered domain, the model should be able to extract spatial-invariant features. This is the main reason why most of the works that were proposed in the scientific literature exploit Convolutional Neural Networks (CNNs). (Tong et al., 2022) developed a CNN model that identifies TC fingerprints according to Satellite Cloud Images of more than 200 TCs over the Northwest Pacific basin, which shows the highest frequency of formation. The results they obtained show that as the TC intensity strengthens, the accuracy improves. In the work of (Nair et al., 2022) TCs are detected from high-resolution satellite images through a multistaged deep learning framework that only considers the shape of the clouds in the images alongside the maximum sustained surface wind speeds data produced by the Joint Typhoon Warning Center (*JTWC*). The framework is made up of a state-of-the-art mask region-CNN (*R-CNN*) detector, a wind speed filter, and a CNN classifier. Segmentation is also added to each detection, fostering the analysis according to the shape and the size of the TC. (Matsuoka et al., 2018) identified TCs and their precursors through an ensemble CNN classifier on twenty-year simulated outgoing longwave radiation simulated by a cloud-resolving global atmospheric model. (Gardoll and Boucher, 2022) adapted and tested a CNN to classify the presence or absence of the cyclonic phenomena in reanalysis maps. The main goal was the evaluation of the performance and sensitivity of the CNN with respect to the learning dataset: indeed, *ERA5* and *MERRA-2* reanalysis datasets were employed. (Pang et al., 2021) propose a detection framework combining the Deep Convolutional Generative Adversarial Networks (*DCGAN*) and You Only Look Once (*YOLO*) v3 models. In particular, the DCGAN is used to augment the dataset by generating synthetic images containing TCs which are employed to pre-train the YOLOv3 model. Then the transfer learning method is applied, by training the detection model with real images and starting from the weights obtained in the pre-training phase.

Another important area related to climate change effects on the environment concerns wildfire risk assessment and prediction. According to researchers, global warming leads to several interdependent effects on both flora and fauna that increases fire risk (Barbero et al., 2015), [19] [20] [21].
The literature accounts for many ML approaches that deal with the Fire Weather Index (*FWI*) prediction in order to use it as a proxy for fire risk assessment. Common ML approaches to fire risk prediction can be divided into two macro categories: 1) burned area extension estimation cases and 2) FWI risk classification works. In regard to the first case, (Omar et al., 2021) proposed a Long Short Term Memory algorithm (*LSTM*)

[19] https://www.bbc.com/news/science-environment-60483431
[20] https://www.nytimes.com/2022/02/23/climate/climate-change-un-wildfire-report.html
[21] https://www.theguardian.com/environment/2022/feb/23/climate-crisis-driving-increase-in-wildfires-across-globe-says-report-aoe

to forecast the total burned area (in hectares, h_a) in the region of Montesano Park by taking in input 12 different features. Among these, they considered the FWI indices as well as other climate drivers. (Safi and Bouroumi, 2013) proposed a work very close to the previous one. As (Omar et al., 2021), they provided a Multi-Layer Perceptron (*MLP*) architecture that takes in input 12 features related to both the FWI and weather conditions.
With regard to FWI risk classification, works like (Kosović and Škurić Kuraži, 2021) and (Rosadi et al., 2022) are extremely meaningful. More in detail, (Kosović and Škurić Kuraži, 2021) develop several state-of-the-art ML algorithms in order to classify FWI risk, such as Random Forest (*RF*), Logistic Regression (*LR*), MLP and others. They make use of an 8-feature vector of climate variables to predict FWI value and therefore its corresponding class. At last, (Rosadi et al., 2022) focused on the estimation of the burned area that could be burnt if nothing is done to prevent fire spread. The novelty of their approach is the use of 12 different climate drivers to develop a Fuzzy C-Means (*FCM*) algorithm.

Physics Informed Machine Learning
Physics Informed Machine Learning (*PIML*) is a cutting-edge research branch that has been developing in recent years. The idea under which PIML is founded is to link ML to the physical world in order to provide physically consistent results. This research topic is fundamental for climate change applications, since PIML models are able to provide physically consistent results, avoiding the uncertainty related to classical Neural Networks algorithms. The literature reviews several efficient ways to embed physical knowledge into a ML model, such as the definition of physics-informed layers, custom losses as well as activation functions (Kashinath et al., 2021).

Among PIML approaches for climate change applications, (Manepalli et al., 2019), developed a Physics Informed (*PI*) Conditional Generative Adversarial Network (*cGAN*). In particular, they propose a novel method for the generation of Snow Water Equivalent (*SWE*) maps. They proposed a modified version of the *Pix2Pix model* (Isola et al., 2017) that was able to embed physical constraints into the model, leading to physically accurate and reliable synthetic SWE samples. In order to enhance physical consistency into the model, in (Manepalli et al., 2019), they propose a modified version of the loss, a PI loss, which embeds domain information as well as SWE characteristics. Indeed, they have formulated three different loss contributions: i) SWE in the sea surface can be only 0, so the loss function heavily affects values greater than 0 along the sea; ii) mountain areas typically have higher values of SWE, hence the network gives more importance to errors in these areas; iii) there is a penalisation of the difference between synthetic SWE and actual one.

Another relevant work in this research field is (Lütjens et al., 2020), which proposes a PI implementation of the Pix2Pix network to assist rescuers during floods. Highly specific evaluation metrics were developed for this task.

4.2 2nd Workshop on Emerging technologies for weather and climate modelling

Following the success of the first workshop organised in 2020 (Aloisio et al., 2020), a second workshop on Emerging technologies for weather and climate was organised to gather scientists and industry experts working on exascale hardware, programming models and machine learning together with the climate/weather community. The main goal of the workshop was to provide an update on the latest developments and early results in areas relevant to the weather and climate community, in particular, for supporting future very high-resolution weather and climate models[22]. This section reports on the workshop organisation and the talks given.

[22] `https://www.esiwace.eu/events/2nd-virtual-workshop-emerging-technologies`

4.2.1 Workshop organisation

The program committee of the workshop was in charge of the organisation and selection of speakers and was composed by:

- Giovanni Aloisio (CMCC)
- Italo Epicoco (CMCC)
- Peter Dueben (ECMWF)
- Rupert Ford (STFC)
- Erwan Raffin (ATOS)
- Donatello Elia (CMCC)

As part of the ESiWACE project office, Dela Spikerman (DKRZ) and Julia Duras (DKRZ) provided support for technical and dissemination aspects.

The workshop was well received by the community with 132 people registered from 22 countries all over the world. The most represented countries were Germany, the UK, Italy, France and USA. The participants were mainly from academia and research centres, although a fair amount of industry representatives also registered (e.g., ARM, ATOS, Google, Intel, Microsoft, NVIDIA, etc.). The level of experience of the participants equally distributed among early-career, mid-level and senior. A news item was published on the ESiWACE portal after completing the workshop[23].

A total of 14 talks were grouped in three topical tracks related to emerging technologies and solutions for climate and weather. The talks were delivered live in a videoconference room and were recorded for later publication on the ESiWACE YouTube channel[24]. Each speakers was assigned with a 20 minutes time slot including question time. Among the talks, 10 were given by representatives of academic institutions and research centres, while 4 were given by industry representatives (the full agenda is available in the Appendix 1 of this document and on the Indico event webpage[25]). The three tracks were:

- *European exascale hardware*
- *Programming models and hardware interplay*
- *Machine learning*

The following subsections summarise the talks in each of the three topical tracks. For each talk, the abstract provided by the speaker is included.

4.2.2 European exascale hardware

The session on European exascale hardware was chaired by Erwan Raffin (ATOS) and provided a snapshot of the current state of the main European pre-exascale machines procured under the *EuroHPC Joint Undertaking*[26], as well as some key projects focusing on exascale systems (i.e., *EUPEX* and *EUPILOT* projects). From the presentation of the pre-exascale European systems it emerged that acceleration computing hardware has become an important part of the infrastructure and that energy efficiency has been a key factor in the design of the systems, given their massive scales. The two projects presented in the session highlighted some of the

[23] https://www.esiwace.eu/news/news/emtech-ws2022-completed
[24] https://www.youtube.com/@esiwace880/
[25] https://indico.dkrz.de/event/45/
[26] https://eurohpc-ju.europa.eu/about/our-supercomputers_en

efforts in building European exa-scale systems and hardware architectures.

Talk1: Leonardo supercomputing system and the national Data Valley action - Sanzio Bassini (CINECA, Italy)

Abstract: *The Leonardo supercomputing system, one of the pre-exascale systems procured by EuroHPC JU, as well as being part of the European supercomputing ecosystem, is integrated in the context of the national Data Valley project, currently being built at the Bologna tecnopolo, where the system Leonardo will be hosted in the new CINECA exascale data center. The brief communication intends to share the state of the art of this action as well as provide the technological details of the supercomputing system and the roadmap for rapid start-up in production of the system.*

Talk2: LUMI - The pre-exascale system in the North - Jenni Kontkanen (CSC – IT Center for Science, Finland)

Abstract: *The EuroHPC initiative is a joint undertaking by the European Commission and 31 countries to establish a world-leading supercomputing and data infrastructure in Europe. One of its first efforts is to install three pre-exascale supercomputers in Europe. I will present one of these systems, LUMI, located in Kajaani, Finland. LUMI is jointly funded and operated by the European commission and a consortium of 10 countries. LUMI is currently the fastest supercomputer in Europe and one of the most powerful and advanced computing systems in the world. In this talk, I will introduce the technical architecture of the LUMI infrastructure, provide a status update of the program, and discuss the opportunities that LUMI provides for climate research.*

Talk3: MareNostrum5 and its Data Center - Sergi Girona (BSC, Spain)

Abstract: *In 2019, the EuroHPC selected the Barcelona Supercomputing Center as host of one of the largest supercomputers in Europe, MareNostrum 5; this new MareNostrum will be a pre-exascale machine, with a peak performance of more than 200 petaflops, 17 times higher than the present MareNostrum 4 and 10,000 times superior to the supercomputer that initiated the saga in 2004: MareNostrum 1. The talk will be about how a data center with the necessary features to host an infrastructure of such magnitude is being prepared -with the problems and solutions involved-, a magnitude that exceeds the current capacities of Torre Girona Chapel. It will also be about the features of the future MareNostrum 5 and its development plan.*

Talk4: EUPEX (European Pilot for Exascale) project on the road to Exascale - Etienne Walter (Atos, France)

Abstract: *The talk will present the objectives of the EUPEX project, and its main purpose: expose at pilot scale the technologies issued from the European Processor Initiative (EPI) project.*

Talk5: EuPILOT - Carlos Puchol (BSC, Spain)

Abstract: *The EuPILOT project aims to build an end-to-end demonstrator of accelerators that could be used in a pre-exascale system, making full use of European and open-source technologies and standards. The project will produce three chip tapeouts. The first will be a test chip to validate the use of the 12nm technology node. The second and third, developed concurrently, will contain a vector accelerator with up to 16 cores and a machine learning and stencil accelerator with up to eight cores, respectively. These will be mounted in modules with LPDDR memory chips. The modules will be installed into accelerator boards going into systems and, when paired with host servers, deployed into liquid immersion tanks with ultra-efficient* *power densities.*

4.2.3 Programming models and hardware interplay

This session was chaired by Andrew Porter (STFC) and described some experiences in the use of Domain Specific Languages (DSLs) and high-level programming models for code portability on heterogeneous infrastructures, as well as opportunities in *DestinE*. In particular, the talks showed some of the experience from both the EU and US sides in using higher-level languages for supporting performance portability of climate/weather modelling systems across different computing devices.

Talk1: PSyclone for LFRic - Iva Kavcic (Met Office, UK)

Abstract: *LFRic is the new weather and climate modelling system being developed by the UK Met Office to replace the existing Unified Model in preparation for exascale computing in the 2020s. LFRic uses the GungHo dynamical core and runs on a semi-structured cubed-sphere mesh. PSyclone is a domain-specific compiler and source-to-source translator developed for use in finite element, finite volume and finite difference codes used in Earth System Models. In essence, it uses domain-specific knowledge to construct a representation of the target code and allows the HPC expert to transform this representation to add various forms of parallelisation (distributed memory, OpenMP threading, OpenMP offload, OpenACC). This talk will give an overview on how PSyclone is used in LFRic, as well as recent examples of the impact of PSyclone on the LFRic performance.*

Talk2: ESiWACE2 DSLs for ICON and NEMO - Carlos Osuna (MeteoSwiss, Switzerland)

Abstract: *One of the objectives of ESiWACE2 is to establish domain-specific languages (DSLs) for community models and evaluate their potential. ESiWACE2 is implementing the use of two main DSLs into several weather and climate models. The PSyclone DSL is used for the development of the LFric model and NEMO while the gt4py DSL is used to re-write the dynamical core of the ICON model. We will present the status of this work and give insights on the main characteristics and key advantages of using these two DSLs: usability, performance portability and long term maintainability.*

Talk3: Experiences with Kokkos in E3SM - Luca Bertagna (Sandia National Lab, US)

Abstract: *The Energy Exascale Earth System Model (E3SM) is a climate model developed by the US Department of Energy. It includes separate components for modeling different parts of earth's climate, including land processes, atmosphere, rivers, ocean, and more, which can all run at different resolutions. In recent years, there has been a push for higher resolutions, which has fueled a large amount of research and code development. The goal is to have a single code base that can perform well across a variery of HPC architectures, a goal usually referred to as performance portability. E3SM has explored several approaches to achieve performance portability, one of which relies on using Kokkos for on-node parallelism. Kokkos is a C++ library that implements a single API that allows developers to efficiently utilize several threading backends, on a variety of HPC architectures. In this presentation, we will discuss our experience in adapting an existing large Fortran code base such as E3SM to use C++, with all its advantages and disadvantages, using Kokkos.*

Talk4: DestinE: opportunities & challenges for digital twins of the Earth System - Balthasar Reuter (ECMWF, International)

Abstract: *Destination Earth is an ambitious initiative by the European Commission to develop a digital twin of Earth on a global scale. In the first phase, ECMWF is commited to delivering the Digital Twin Engine and* *the first two digital twins on weather-induced extremes and climate change adaptation. The computational resources available from the EuroHPC Joint Undertaking provide opportunity for simulations at scale and great detail. However, sustainable adaptation of ECMWF's Integrated Forecasting System to the GPU-accelerated supercomputers alongside scientific developments poses a major challenge. We present the principles underpinning this effort and early results obtained for some components of the IFS.*

4.2.4 Machine learning

The last session of the workshop was chaired by Peter Dueben (ECMWF) and illustrated some state-of-the-art work in the field of Machine Learning for weather and climate applications, related, for example, to model augmentation and replacement. From the session it appeared clear that there are several research efforts going on (from both academy and industry) in using data-driven approaches in Earth system modelling (e.g., combining physical models and ML models), as well as various challenges that still need to be addressed.

Talk1: Atmospheric Physics-Guided Machine Learning for Climate Modeling and Weather Forecasting - Tom Beucler (EPSL, Switzerland)

Abstract: *Data-driven algorithms, in particular neural networks, can (1) emulate the effect of unresolved processes in coarse-resolution climate models if trained on high-resolution simulation data; and (2) significantly improve the skill of numerical weather forecasts at low computational cost if trained on observations. However, they may violate key physical constraints and make large errors when evaluated outside of their training set. I will share progress towards overcoming these two challenges in the case of machine learning (1) the effect of subgrid-scale convection and clouds on the large-scale climate; and (2) the bias correction of near-surface temperature and humidity to post-process numerical weather predictions. First, physical constraints can be enforced in neural networks, either approximately by adapting the loss function or to within machine precision by adapting the architecture. Second, as these physical constraints are insufficient to guarantee generalizability, we additionally propose to transform the inputs and outputs of machine learning algorithms using established physical rescalings to help them generalize to unseen climates and locations. Overall, these results suggest that explicitly incorporating physical knowledge into data-driven models for weather and climate applications may improve their consistency, stability, and ability to generalize across atmospheric regimes.*

Talk2: A Generative Deep Learning Approach to Stochastic Downscaling of Precipitation Forecasts - Lucy Harris (Oxford, UK)

Abstract: *Several key processes affecting precipitation distribution and intensity occur below the resolved scale of global weather models. Generative adversarial networks (GANs) have been demonstrated by the computer vision community to be successful at super-resolution problems, i.e., learning to add fine-scale structure to coarse images. We demonstrate this approach can be extended to the more challenging problem of increasing the accuracy and resolution of comparatively low-resolution input from a weather forecasting model. We test our models and show that they perform in a range of scenarios, including heavy rainfall.*

Talk3: Building Digital Twins of the Earth for NVIDIA's Earth-2 Initiative - Karthik Kashinath and Mike Pritchard (NVIDIA, US)

Abstract: *NVIDIA is committed to helping address climate change. Recently our CEO announced the Earth-2 initiative, which aims to build digital twins of the Earth and a dedicated supercomputer, E-2, to power them. Two central goals of this initiative are: (i) Computational: Enable high-resolution climate predictions with credible cloud physics; and (ii) Societal: Nimbly serve interactive, useful, next-generation climate predictions* *via NVIDIA Omniverse. These predictions will help plan for the disastrous impacts of climate change well in advance and develop strategies to mitigate and adapt to change. Building ML-driven digital twins of the Earth requires reliable ground truth training data. Today's storm-resolving climate models at km-scale are the most credible ground truth one can use for future climate change. However, these models are prohibitively expensive to run and produce overwhelming amounts of data. Further, building a robust and generalizable digital twin requires training data from hundreds of diverse trajectories to sample the long tails of extreme climate change. We present our four-part strategy towards building digital twins: (i) HOP: 100X speedup of km-scale hybrid ML-climate model using ML-physics coarse-graining; (ii) SKIP: 10,000X speedup of sub-km-scale multi-scale hybrid ML-climate model for explicit low cloud dynamics; (iii) TUNE: Reinforcement learning-based auto-calibration of km-scale climate models to validate and improve sub-grid uncertainty estimates; (iv) LEAP: 10,000X speedup via tethering an ML surrogate to checkpoints of km-scale accelerated hybrid ML-climate models. We discuss the implications of the computational speedup and data compression factors in each of the four parts for building digital twins of the Earth. We present recent developments of our ML surrogate, FourCastNet, including advances in scale (of ML model, compute, and size of input training state vector), uncertainty calibration, high-resolution regionalization, and a real-time weather prediction system. We conclude with a roadmap of Earth-2 that encompasses weather and climate goals and outlines engineering innovations required for the breakthroughs that building digital twins of the Earth demands.*

Talk4: Deep learning and differentiable simulations - Stephan Hoyer (Google, US)

Abstract: *How could deep learning be transformative for weather and climate simulation? In this talk, I'll give an overview of a line of research at Google, on how deep learning can improve numerical solvers via end-to-end optimization with differentiable simulators. The approach allows for both increased accuracy and higher performance on hardware accelerators, within the structure of traditional numerical models. I will share our results on 2D turbulence and our progress on building a new global atmospheric model.*

Talk5: Machine learning for weather forecasting: successes, challenges, and the future - Jonathan Weyn (Microsoft, US)

Abstract: *The last few years have seen massive developments in machine learning for weather forecasting, from forecast post-processing to emulating numerical weather models to fully operational machine learning systems. In this talk, I will cover a touch upon a few of these topics, including NWP replacement, ensemble post-processing with transformers, and the operational precipitation nowcasting at Microsoft Weather. This discussion will detail some of the successes, challenges, and future directions of machine learning in weather.*

5 References

Gabriele Accarino, Marco Chiarelli, Francesco Immorlano, Valeria Aloisi, Andrea Gatto, and Giovanni Aloisio. Msg-gan-sd: A multi-scale gradients gan for statistical downscaling of 2-meter temperature over the euro-cordex domain. *AI*, 2(4):600–620, 2021. ISSN 2673-2688. doi: 10.3390/ai2040036. URL https://www.mdpi.com/2673-2688/2/4/36.

S.V. Adams, R.W. Ford, M. Hambley, J.M. Hobson, I. Kavčič, C.M. Maynard, T. Melvin, E.H. Müller, S. Mullerworth, A.R. Porter, M. Rezny, B.J. Shipway, and R. Wong. Lfric: Meeting the challenges of scalability and performance portability in weather and climate models. *Journal of Parallel and Distributed Computing*, 132:383–396, 2019. ISSN 0743-7315. doi: https://doi.org/10.1016/j.jpdc.2019.02.007. URL https://www.sciencedirect.com/science/article/pii/S0743731518305306.

S. H. Alemohammad, J. Kolassa, C. Prigent, F. Aires, and P. Gentine. Global downscaling of remotely sensed soil moisture using neural networks. *Hydrology and Earth System Sciences*, 22(10):5341–5356, 2018. doi: 10.5194/hess-22-5341-2018. URL https://hess.copernicus.org/articles/22/5341/2018/.

Moteb Alghamdi, Graham Riley, and Mike Ashworth. A comparison of vivado hls, sdsoc c++ and opencl for porting a matrix-vector-based climate model mini-app to fpgas. In *PDPTA'21-The 27th Int'l Conference on Parallel and Distributed Processing Techniques and Applications*, 2021. [Accepted/In Press].

Giovanni Aloisio, Graham Riley, Sandro Fiore, and Carlos Osuna. First white paper on community guidelines on the use, value and applicability of emerging technologies in climate and weather applications, August 2020. URL https://doi.org/10.5281/zenodo.4001485.

Mike Ashworth, Graham D Riley, Andrew Attwood, and John Mawer. First steps in porting the lfric weather and climate model to the fpgas of the euroexa architecture. *Scientific Programming*, 2019, 2019. doi: 10.1155/2019/7807860.

J. Baño Medina, R. Manzanas, and J. M. Gutiérrez. Configuration and intercomparison of deep learning neural models for statistical downscaling. *Geoscientific Model Development*, 13(4):2109–2124, 2020. doi: 10.5194/gmd-13-2109-2020. URL https://gmd.copernicus.org/articles/13/2109/2020/.

Jorge Baño-Medina, José Manuel Gutiérrez, and Sixto Herrera. Deep neural networks for statistical downscaling of climate change projections 1. 2018. URL https://www.semanticscholar.org/paper/Deep-Neural-Networks-for-Statistical-Downscaling-of-Ba%C3%B1o-Medina-Guti%C3%A9rrez/169752689db16b6e5e09c532a855405b301fd452.

R. Barbero, J. T. Abatzoglou, N. K. Larkin, C. A. Kolden, and B. Stocks. Climate change presents increased potential for very large fires in the contiguous united states. *International Journal of Wildland Fire*, 24(7):892, 2015. doi: 10.1071/WF15083. URL https://doi.org/10.1071/WF15083.

Paulina Bartkowiak, Mariapina Castelli, and Claudia Notarnicola. Downscaling land surface temperature from modis dataset with random forest approach over alpine vegetated areas. *Remote Sensing*, 11(11), 2019. ISSN 2072-4292. doi: 10.3390/rs11111319. URL https://www.mdpi.com/2072-4292/11/11/1319.

David A Beckingsale, Jason Burmark, Rich Hornung, Holger Jones, William Killian, Adam J Kunen, Olga Pearce, Peter Robinson, Brian S Ryujin, and Thomas RW Scogland. Raja: Portable performance for large-scale scientific applications. In *2019 ieee/acm international workshop on performance, portability and productivity in hpc (p3hpc)*, pages 71–81. IEEE, 2019. doi: 10.1109/P3HPC49587.2019.00012.

Tal Ben-Nun, Linus Groner, Florian Deconinck, Tobias Wicky, Eddie Davis, Johann Dahm, Oliver D. Elbert, Rhea George, Jeremy McGibbon, Lukas Trümper, Elynn Wu, Oliver Fuhrer, Thomas Schulthess, and Torsten Hoefler. Productive performance engineering for weather and climate modeling with python, 2022. URL https://arxiv.org/abs/2205.04148.

Luca Bertagna, Michael Deakin, Oksana Guba, Daniel Sunderland, Andrew M Bradley, Irina K Tezaur, Mark A Taylor, and Andrew G Salinger. Hommexx 1.0: a performance-portable atmospheric dynamical core for the energy exascale earth system model. *Geoscientific Model Development*, 12(4):1423–1441, 2019. doi: 10.5194/gmd-12-1423-2019. URL `https://gmd.copernicus.org/articles/12/1423/2019/`.

Yi-Chia Chang, Ralph Acierto, Tomoaki Itaya, Kawasaki Akiyuki, and Ching-Pin Tung. A Deep Learning Approach to Downscaling Precipitation and Temperature over Myanmar. In *EGU General Assembly Conference Abstracts*, EGU General Assembly Conference Abstracts, page 4120, April 2018.

Valentin Clement, Sylvaine Ferrachat, Oliver Fuhrer, Xavier Lapillonne, Carlos E. Osuna, Robert Pincus, Jon Rood, and William Sawyer. The claw dsl: Abstractions for performance portable weather and climate models. In *Proceedings of the Platform for Advanced Scientific Computing Conference*, PASC '18, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450358910. doi: 10.1145/3218176.3218226. URL `https://doi.org/10.1145/3218176.3218226`.

Tom Deakin, Simon McIntosh-Smith, James Price, Andrei Poenaru, Patrick Atkinson, Codrin Popa, and Justin Salmon. Performance portability across diverse computer architectures. In *2019 IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC (P3HPC)*, pages 1–13, 2019. doi: 10.1109/P3HPC49587.2019.00006.

Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *Computer Vision – ECCV 2014*, pages 184–199, Cham, 2014. Springer International Publishing. ISBN 978-3-319-10593-2. doi: https://doi.org/10.1007/978-3-319-10593-2_13.

Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(2):295–307, 2016. doi: 10.1109/TPAMI.2015.2439281.

Jack Dongarra, Steven Gottlieb, and William T. C. Kramer. Race to exascale. *Computing in Science & Engineering*, 21(1):4–5, 2019. doi: 10.1109/MCSE.2018.2882574.

Darshana Duhan and Ashish Pandey. Statistical downscaling of temperature using three techniques in the tons river basin in central india. *Theoretical and applied climatology*, 121(3):605–622, 2015. doi: https://doi.org/10.1007/s00704-014-1253-5.

Yves Durand, Paul M. Carpenter, Stefano Adami, Angelos Bilas, Denis Dutoit, Alexis Farcy, Georgi Gaydadjiev, John Goodacre, Manolis Katevenis, Manolis Marazakis, Emil Matus, Iakovos Mavroidis, and John Thomson. Euroserver: Energy efficient node for european micro-servers. In *17th Euromicro Conference on Digital System Design*, pages 206–213, 2014. doi: 10.1109/DSD.2014.15.

H Carter Edwards, Christian R Trott, and Daniel Sunderland. Kokkos: Enabling manycore performance portability through polymorphic memory access patterns. *Journal of parallel and distributed computing*, 74(12):3202–3216, 2014. doi: https://doi.org/10.1016/j.jpdc.2014.07.003. URL `https://www.sciencedirect.com/science/article/pii/S0743731514001257`.

Kerry Emanuel. Tropical cyclones. *Annual Review of Earth and Planetary Sciences*, 31(1):75–104, 2003. doi: 10.1146/annurev.earth.31.100901.141259. URL `https://doi.org/10.1146/annurev.earth.31.100901.141259`.

Wei Fang, Qiongying Xue, Liang Shen, and Victor S. Sheng. Survey on the application of deep learning in extreme weather prediction. *Atmosphere*, 12(6), 2021. ISSN 2073-4433. doi: 10.3390/atmos12060661. URL `https://www.mdpi.com/2073-4433/12/6/661`.

Lu Gao, Karsten Schulz, and Matthias Bernhardt. Statistical downscaling of era-interim forecast precipitation data in complex terrain using lasso algorithm. *Advances in Meteorology*, 2014, 2014. doi: https://doi.org/10.1155/2014/472741.

S. Gardoll and O. Boucher. Classification of tropical cyclone containing images using a convolutional neural network: performance and sensitivity to the learning dataset. *EGUsphere*, 2022:1–29, 2022. doi: 10.5194/egusphere-2022-147. URL https://egusphere.copernicus.org/preprints/egusphere-2022-147/.

Aneesh Goly, Ramesh S. V. Teegavarapu, and Arpita Mondal. Development and evaluation of statistical downscaling models for monthly precipitation. *Earth Interactions*, 18(18):1 – 28, 2014. doi: 10.1175/EI-D-14-0024.1. URL https://journals.ametsoc.org/view/journals/eint/18/18/ei-d-14-0024.1.xml.

Brian Groenke, Luke Madaus, and Claire Monteleoni. Climalign: Unsupervised statistical downscaling of climate variables via normalizing flows. In *Proceedings of the 10th International Conference on Climate Informatics*, CI2020, page 60–66, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450388481. doi: 10.1145/3429309.3429318. URL https://doi.org/10.1145/3429309.3429318.

Lucy Harris, Andrew T. T. McRae, Matthew Chantry, Peter D. Dueben, and Tim N. Palmer. A generative deep learning approach to stochastic downscaling of precipitation forecasts. *Journal of Advances in Modeling Earth Systems*, 14(10):e2022MS003120, 2022. doi: https://doi.org/10.1029/2022MS003120. URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2022MS003120.

Hans Hersbach, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, Joaquín Muñoz-Sabater, Julien Nicolas, Carole Peubey, Raluca Radu, Dinand Schepers, Adrian Simmons, Cornel Soci, Saleh Abdalla, Xavier Abellan, Gianpaolo Balsamo, Peter Bechtold, Gionata Biavati, Jean Bidlot, Massimo Bonavita, Giovanna De Chiara, Per Dahlgren, Dick Dee, Michail Diamantakis, Rossana Dragani, Johannes Flemming, Richard Forbes, Manuel Fuentes, Alan Geer, Leo Haimberger, Sean Healy, Robin J. Hogan, Elías Hólm, Marta Janisková, Sarah Keeley, Patrick Laloyaux, Philippe Lopez, Cristina Lupu, Gabor Radnoti, Patricia de Rosnay, Iryna Rozum, Freja Vamborg, Sebastien Villaume, and Jean-Noël Thépaut. The era5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049, 2020. doi: https://doi.org/10.1002/qj.3803. URL https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.3803.

Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, nov 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735. URL https://doi.org/10.1162/neco.1997.9.8.1735.

X. Huang. Deep-learning based climate downscaling using the super-resolution method: a case study over the western us. *Geoscientific Model Development Discussions*, 2020:1–18, 2020. doi: 10.5194/gmd-2020-214. URL https://gmd.copernicus.org/preprints/gmd-2020-214/.

Kevin Höhlein, Michael Kern, Timothy Hewson, and Rüdiger Westermann. A comparative study of convolutional neural network models for wind field downscaling. *Meteorological Applications*, 27(6):e1961, 2020. doi: https://doi.org/10.1002/met.1961. URL https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/met.1961.

Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei Efros. Image-to-image translation with conditional adversarial networks. pages 5967–5976, 07 2017. doi: 10.1109/CVPR.2017.632.

Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution, 2016. URL https://arxiv.org/abs/1603.08155.

Dr C. S. KanimozhiSelvi and G. Sowmiya. Prediction of extreme weather events using machine learning technique. *International journal of applied engineering research*, 14(4):924–929, 2019. ISSN 0973-4562. URL https://www.ripublication.com/ijaer19/ijaerv14n4_12.pdf.

K. Kashinath, M. Mustafa, A. Albert, J-L. Wu, C. Jiang, S. Esmaeilzadeh, K. Azizzadenesheli, R. Wang, A. Chattopadhyay, A. Singh, A. Manepalli, D. Chirila, R. Yu, R. Walters, B. White, H. Xiao, H. A. Tchelepi, P. Marcus, A. Anandkumar, P. Hassanzadeh, and null Prabhat. Physics-informed machine learning: case studies for weather and climate modelling. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 379(2194):20200093, 2021. doi: 10.1098/rsta.2020.0093. URL https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2020.0093.

Michael Kern, Kevin Höhlein, Timothy Hewson, and Rüdiger Westermann. Towards Operational Downscaling of Low Resolution Wind Fields using Neural Networks. In *EGU General Assembly Conference Abstracts*, EGU General Assembly Conference Abstracts, page 5447, May 2020. doi: 10.5194/egusphere-egu2020-5447.

Kawthar Shafie Khorassani, Ching-Hsiang Chu, Hari Subramoni, and Dhabaleswar K Panda. Performance evaluation of mpi libraries on gpu-enabled openpower architectures: Early experiences. In *International Conference on High Performance Computing*, pages 361–378. Springer, 2019. doi: https://doi.org/10.1007/978-3-030-34356-9_28.

Heewon Kim, Myungsub Choi, Bee Lim, and Kyoung Mu Lee. Task-aware image downscaling. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV 2018*, pages 419–434, Cham, 2018. Springer International Publishing. ISBN 978-3-030-01225-0. doi: https://doi.org/10.1007/978-3-030-01225-0_25.

Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1646–1654, 2016. doi: 10.1109/CVPR.2016.182.

Volodymyr Kindratenko and Pedro Trancoso. Trends in high-performance computing. *Computing in Science & Engineering*, 13(3):92–95, 2011. doi: 10.1109/MCSE.2011.52.

Ivana Nižetić Kosović and Diana Škurić Kuraži. Machine learning approach in fire risk estimation. In *2021 International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, pages 1–6, 2021. doi: 10.23919/SoftCOM52868.2021.9559064.

Jussi Leinonen, Daniele Nerini, and Alexis Berne. Stochastic super-resolution for downscaling time-evolving atmospheric fields with a generative adversarial network. *IEEE Transactions on Geoscience and Remote Sensing*, 59(9):7211–7223, sep 2021. doi: 10.1109/tgrs.2020.3032790. URL https://doi.org/10.1109%2Ftgrs.2020.3032790.

Xinyi Li, Zhong Li, Wendy Huang, and Pengxiao Zhou. Performance of statistical and machine learning ensembles for daily temperature downscaling. *Theoretical and Applied Climatology*, 140(1):571–588, 2020. doi: https://doi.org/10.1007/s00704-020-03098-3.

Yangxiaoyue Liu, Yaping Yang, Wenlong Jing, and Xiafang Yue. Comparison of different machine learning approaches for monthly satellite-based soil moisture downscaling over northeast china. *Remote Sensing*, 10 (1), 2018. ISSN 2072-4292. doi: 10.3390/rs10010031. URL https://www.mdpi.com/2072-4292/10/1/31.

Björn Lütjens, Brandon Leshchinskiy, Christian Requena-Mesa, Farrukh Chishtie, Natalia Diaz Rodriguez, Océane Boulais, Aaron Piña, Dava Newman, Alexander Lavin, Yarin Gal, and Chedy Raïssi. Physics-informed gans for coastal flood visualization. 10 2020. doi: https://doi.org/10.48550/arXiv.2010.08103.

Ashray Manepalli, Adrian Albert, Alan Rhoades, and Daniel Feldman. Emulating numeric hydroclimate models with physics-informed cgans. In *NeurIPS 2019 Workshop on Tackling Climate Change with Machine Learning*, 2019. URL https://www.climatechange.ai/papers/neurips2019/39.

Igor L Markov. Limits on fundamental limits to computation. *Nature*, 512(7513):147–154, 2014. doi: https://doi.org/10.1038/nature13570.

Daisuke Matsuoka, Masuo Nakano, Daisuke Sugiyama, and Seiichi Uchida. Deep learning approach for detecting tropical cyclones and their precursors in the simulation by a cloud-resolving global nonhydrostatic atmospheric model. *Progress in Earth and Planetary Science*, 5(1):1–16, 2018. doi: https://doi.org/10.1186/s40645-018-0245-y.

David Mendes and José A. Marengo. Temporal downscaling: a comparison between artificial neural network and autocorrelation techniques over the Amazon Basin in present and future climate change scenarios. *Theoretical and Applied Climatology*, 100(3-4):413–421, May 2010. doi: 10.1007/s00704-009-0193-y.

Qinghua Miao, Baoxiang Pan, Hao Wang, Kuolin Hsu, and Soroosh Sorooshian. Improving monsoon precipitation prediction using combined convolutional and long short term memory neural network. *Water*, 11(5), 2019. ISSN 2073-4441. doi: 10.3390/w11050977. URL https://www.mdpi.com/2073-4441/11/5/977.

Xiaoxiao Min, Ziqiang Ma, Jintao Xu, Kang He, Zhige Wang, Qingliang Huang, and Jun Li. Spatially downscaling imerg at daily scale using machine learning approaches over zhejiang, southeastern china. *Frontiers in Earth Science*, 8, 2020. ISSN 2296-6463. doi: 10.3389/feart.2020.00146. URL https://www.frontiersin.org/articles/10.3389/feart.2020.00146.

Saptarshi Misra, Sudeshna Sarkar, and Pabitra Mitra. Statistical downscaling of precipitation using long short-term memory recurrent neural networks. *Theoretical and applied climatology*, 134(3):1179–1196, 2018. doi: https://doi.org/10.1007/s00704-017-2307-2.

G. E. Moore. Cramming more components onto integrated circuits, reprinted from electronics, volume 38, number 8, april 19, 1965, pp.114 ff. *IEEE Solid-State Circuits Society Newsletter*, 11(3):33–35, 2006. doi: 10.1109/N-SSC.2006.4785860. URL http://doi.org/10.1109/N-SSC.2006.4785860.

Soukayna Mouatadid, Steve Easterbrook, and Andre R. Erler. A machine learning approach to non-uniform spatial downscaling of climate variables. In *2017 IEEE International Conference on Data Mining Workshops (ICDMW)*, pages 332–341, 2017. doi: 10.1109/ICDMW.2017.49.

Aravind Nair, K. S. S. Sai Srujan, Sayali R. Kulkarni, Kshitij Alwadhi, Navya Jain, Hariprasad Kodamana, S. Sandeep, and Viju O. John. A deep learning framework for the detection of tropical cyclones from satellite images. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5, 2022. doi: 10.1109/LGRS.2021.3131638.

Philipp Neumann, Peter Düben, Panagiotis Adamidis, Peter Bauer, Matthias Brück, Luis Kornblueh, Daniel Klocke, Bjorn Stevens, Nils Wedi, and Joachim Biercamp. Assessing the scales in numerical weather and climate predictions: will exascale be the rescue? *Philosophical Transactions of the Royal Society A*, 377 (2142):20180148, 2019. doi: https://doi.org/10.1098/rsta.2018.0148.

Naaman Omar, Adel Al-zebari, and Abdulkadir Sengur. Deep learning approach to predict forest fires using meteorological measurements. In *2021 2nd International Informatics and Software Engineering Conference (IISEC)*, pages 1–4, 2021. doi: 10.1109/IISEC54230.2021.9672446.

Baoxiang Pan, Kuolin Hsu, Amir AghaKouchak, and Soroosh Sorooshian. Improving precipitation estimation using convolutional neural network. *Water Resources Research*, 55(3):2301–2321, 2019. doi: https://doi.org/10.1029/2018WR024090. URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2018WR024090.

Shanchen Pang, Pengfei Xie, Danya Xu, Fan Meng, Xixi Tao, Bowen Li, Ying Li, and Tao Song. Ndftc: A new detection framework of tropical cyclones from meteorological satellite images with deep transfer learning. *Remote Sensing*, 13(9), 2021. ISSN 2072-4292. doi: 10.3390/rs13091860. URL https://www.mdpi.com/2072-4292/13/9/1860.

Dongwon Park, Jisoo Kim, and Se Young Chun. Down-scaling with learned kernels in multi-scale deep neural networks for non-uniform single image deblurring, 2019. URL https://arxiv.org/abs/1903.10157.

Yu Qi. Literature survey of the use of neural networks in earth system modelling. Msc dissertation, University of Mancheste, 2019. URL https://drive.google.com/file/d/1nFu5TOMf2FAEuw1LeBxy_Ej08dmF_zYV/view?usp=sharing.

Daniel A. Reed and Jack Dongarra. Exascale computing and big data. *Commun. ACM*, 58(7):56–68, 6 2015. ISSN 0001-0782. doi: 10.1145/2699414. URL https://doi.org/10.1145/2699414.

Graham D. Riley, Giovanni Aloisio, and Donatello Elia. Machine Learning Workshop: "New Opportunities in ML/AI for Weather and Climate Modelling" - A virtual workshop held jointly with IS-ENES3 (D2.10), May 2021. URL https://doi.org/10.5281/zenodo.4836892.

Dedi Rosadi, Deasy Arisanty, and Dina Agustina. Prediction of forest fire using neural networks with backpropagation learning and exreme learning machine approach using meteorological and weather index variables. *MEDIA STATISTIKA*, 14(2):118–124, 2022. ISSN 2477-0647. doi: 10.14710/medstat.14.2.118-124. URL https://ejournal.undip.ac.id/index.php/media_statistika/article/view/34908.

D. A. Sachindra, F. Huang, A. Barton, and B. J. C. Perera. Least square support vector and multi-linear regression for statistically downscaling general circulation model outputs to catchment streamflows. *International Journal of Climatology*, 33(5):1087–1106, 2013. doi: https://doi.org/10.1002/joc.3493. URL https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/joc.3493.

D.A. Sachindra, K. Ahmed, Md. Mamunur Rashid, S. Shahid, and B.J.C. Perera. Statistical downscaling of precipitation using machine learning techniques. *Atmospheric Research*, 212:240–258, 2018. ISSN 0169-8095. doi: https://doi.org/10.1016/j.atmosres.2018.05.022. URL https://www.sciencedirect.com/science/article/pii/S0169809517310141.

Youssef Safi and A. Bouroumi. Prediction of forest fires using artificial neural networks. *Applied Mathematical Sciences*, 7:271–286, 01 2013. doi: 10.12988/ams.2013.13025.

Amir Hossein Salimi, Jafar Masoompour Samakosh, Ehsan Sharifi, Mohammad Reza Hassanvand, Amir Noori, and Hary von Rautenkranz. Optimized artificial neural networks-based methods for statistical downscaling of gridded precipitation data. *Water*, 11(8), 2019. ISSN 2073-4441. doi: 10.3390/w11081653. URL https://www.mdpi.com/2073-4441/11/8/1653.

Tsuyoshi Thomas Sekiyama. Statistical downscaling of temperature distributions from the synoptic scale to the mesoscale using deep convolutional neural networks. *arXiv preprint arXiv:2007.10839*, 2020. doi: https://doi.org/10.48550/arXiv.2007.10839.

E. Sharifi, B. Saghafian, and R. Steinacker. Downscaling satellite precipitation estimates with multiple linear regression, artificial neural networks, and spline interpolation techniques. *Journal of Geophysical Research: Atmospheres*, 124(2):789–805, 2019. doi: https://doi.org/10.1029/2018JD028795. URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2018JD028795.

Xiaoming Shi. Enabling smart dynamical downscaling of extreme precipitation events with machine learning. *Geophysical Research Letters*, 47(19):e2020GL090309, 2020. doi: https://doi.org/10.1029/2020GL090309. URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2020GL090309. e2020GL090309 10.1029/2020GL090309.

Lei Sun and Yufeng Lan. Statistical downscaling of daily temperature and precipitation over china using deep learning neural models: Localization and comparison with other methods. *International Journal of Climatology*, 41(2):1128–1147, 2021. doi: https://doi.org/10.1002/joc.6769. URL https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/joc.6769.

B. Tong, X. Sun, J. Fu, Y. He, and P. Chan. Identification of tropical cyclones via deep convolutional neural network based on satellite cloud images. *Atmospheric Measurement Techniques*, 15(6):1829–1848, 2022. doi: 10.5194/amt-15-1829-2022. URL https://amt.copernicus.org/articles/15/1829/2022/.

Duong Tran Anh, Song P. Van, Thanh D. Dang, and Long P. Hoang. Downscaling rainfall using deep learning long short-term memory and feedforward neural network. *International Journal of Climatology*, 39 (10):4170–4188, 2019. doi: https://doi.org/10.1002/joc.6066. URL https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/joc.6066.

Thomas Vandal, Evan Kodra, Sangram Ganguly, Andrew Michaelis, Ramakrishna Nemani, and Auroop R. Ganguly. Deepsd: Generating high resolution climate change projections through single image super-resolution. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '17, page 1663–1672, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450348874. doi: 10.1145/3097983.3098004. URL https://doi.org/10.1145/3097983.3098004.

Thomas Vandal, Evan Kodra, Sangram Ganguly, Andrew Michaelis, Ramakrishna Nemani, and Auroop R Ganguly. Generating high resolution climate change projections through single image super-resolution: An abridged version. *International Joint Conferences on Artificial Intelligence Organization*, 2018. doi: 10.24963/ijcai.2018/759. URL https://par.nsf.gov/biblio/10111432.

Thomas Vandal, Evan Kodra, and Auroop R. Ganguly. Intercomparison of machine learning methods for statistical downscaling: the case of daily and extreme precipitation. *Theoretical and Applied Climatology*, 137(1-2):557–570, July 2019. doi: 10.1007/s00704-018-2613-3.

Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In Laura Leal-Taixé and Stefan Roth, editors, *Computer Vision – ECCV 2018 Workshops*, pages 63–79, Cham, 2019. Springer International Publishing. ISBN 978-3-030-11021-5. doi: https://doi.org/10.1007/978-3-030-11021-5_5.

B. L. White, A. Singh, and A. Albert. Downscaling Numerical Weather Models with GANs. In *AGU Fall Meeting Abstracts*, volume 2019, pages GC43D–1357, December 2019.

Ren Xu, Nengcheng Chen, Yumin Chen, and Zeqiang Chen. Downscaling and projection of multi-cmip5 precipitation using machine learning methods in the upper han river basin. *Advances in Meteorology*, Vol. 2020, 2020. doi: https://doi.org/10.1155/2020/8680436.

Xu Yan. Characteristics and psychological coping in the post-disaster era. *Chinese agency logistics*, (6):68–69, 6 2020.

6 Changes made and/or difficulties encountered, if any

This deliverable has been postponed by 3 months, differently from what was originally planned in the DoA, i.e., from August to November 2022, since the associated workshop has also been moved from May to October 2022. This change of planning was meant to avoid overlapping with the *ENES HPC workshop* organised in May in the context of the IS-ENES3 project in collaboration with ESiWACE2[27]. The ENES HPC workshop covered in fact very similar topics and targeted the same community. Moreover, the leadership of the deliverable has been reassigned from UNIMAN to CMCC. It is worth mentioning that UNIMAN contribution to this activity (until April 2021, before leaving the task) has also been included in this document.

7 How this deliverable contributes to the European strategies for HPC

This document reports about some of the latest developments from literature and early work from the community addressing the use of large-scale HPC infrastructures and novel data-driven solutions for climate and weather applications. Moreover, the European Exascale Hardware session of the workshop (see section 4.2.2) presented (i) the pre-exascale system procured under the EuroHPC JU allowing the audience to interact with the scientists and experts of the computing centres to better understand their capabilities and possible usages, as well as (ii) some of the pilot project aiming to increase the European leadership in HPC technologies. Overall, this contributes to strengthening the interaction between the scientist in the climate/weather community and the HPC experts, as well as facilitating knowledge sharing, in line with the goal of preparing the community to effectively use the world-class HPC infrastructure that are currently being built in Europe.

8 Sustainability

This workshop follows the path of the previous workshop organised in the context of ESiWACE2 for bringing together the community working on emerging technologies for weather forecast and climate modelling: the *First Virtual Workshop on Emerging Technologies for Weather and Climate Modelling*[28], June 2020 and the *IS-ENES3/ESiWACE ML/AI in Weather and Climate modelling workshop*[29], March 2021. The outcomes and key points from the workshops were respectively reported in D2.6 (Aloisio et al., 2020) and D2.10 (Riley et al., 2021). Moreover, it extends the inital set of guidelines reported in the first white paper (Aloisio et al., 2020), by providing a wider analysis of the literature and state-of-the-art solutions.

Similarly to the previous workshop, the workshop reported in this document was organised as a virtual event in order to widen the potential audience following the talks, and also allow speakers and audience from different countries worldwide to get together more easily. In particular, in order to strengthen the link with the industry, the workshop involved contributions of an industrial partner from the ESiWACE2 project, i.e., ATOS (for both the organisation of the workshop and a talk), as well as from US companies like NVIDIA, Google and Microsoft.

Finally, the findings and collection of resources (e.g., papers, online articles, video recordings) highlighted in this second white paper can represent a valuable source for scientists from the climate and weather community interested in trending and emerging solutions.

[27] https://www.esiwace.eu/events/7th-enes-hpc-workshop
[28] https://www.esiwace.eu/events/virtual-ws-on-emerging-technologies/virtual-ws-on-emerging-technologies
[29] https://is.enes.org/workshops-detailed/#ML-AI-WS

9 Dissemination, Engagement and Uptake of Results

9.1 Target audience

As indicated in the Description of the Action, the audience for this deliverable is:

X	The general public (PU)
	The project partners, including the Commission services (PP)
	A group specified by the consortium, including the Commission services (RE)
	This reports is confidential, only for members of the consortium, including the Commission services (CO)

We ensure the uptake of the deliverable by the targeted audience through the following actions

The workshop associated with this deliverable has been widely announced on community-related mailing lists, social networks and websites[30]. The slides of the presentations have been made available for download for anyone interested in the workshop[31]. Moreover, to further reach out towards the community the recordings from the workshop talks have been published on the ESiWACE YouTube channel[32]. The availability of the workshop material has also been announced on mailing lists, social channels and on the ESiWACE project website[33].
Furthermore, this deliverable will be disseminated via the ESiWACE website, in the ESiWACE newsletter and via Twitter, as soon as it is published on Zenodo.

9.2 Record of dissemination/engagement activities linked to this deliverable

See Table 1.

Table 1: Record of dissemination / engagement activities linked to this deliverable

Type of dissemination and communication activities	Details	Date and location of the event	Type of audience	Zenodo Link	Estimated number of persons reached
Organisation of a workshop	ESiWACE2 Second Virtual Workshop on Emerging Technologies for Weather and Climate Modelling	7 October 2022, Online	Science and industry	N.A.	132

[30] https://www.esiwace.eu/news/news/emtech-ws2022
[31] https://indico.dkrz.de/event/45/
[32] https://www.youtube.com/@esiwace880/
[33] https://www.esiwace.eu/news/news/emtech-ws2022-completed

9.3 Publications in preparation OR submitted

See Table 2.

Table 2: Publications related to this deliverable

In preparation OR submitted?	Title	All authors	Title of the periodical or the series	Is/Will open access be provided to this publication?
Submitted	Machine Learning Emulation of 3D Cloud Radiative Effects	Meyer, David and Hogan, Robin J. and Dueben, Peter D. and Mason, Shannon L.	Journal of Advances in Modeling Earth Systems	Yes
Submitted	Machine Learning Emulation of Gravity Wave Drag in Numerical Weather Forecasting	Chantry, Matthew and Hatfield, Sam and Dueben, Peter and Polichtchouk, Inna and Palmer, Tim	Journal of Advances in Modeling Earth Systems	Yes
Submitted	A Generative Deep Learning Approach to Stochastic Downscaling of Precipitation Forecasts	Harris, Lucy and McRae, Andrew T. T. and Chantry, Matthew and Dueben, Peter D. and Palmer, Tim N.	Journal of Advances in Modeling Earth Systems	Yes
Submitted	Comparison of Vivado HLS, SDSoC C++ and OpenCL for Porting a Matrix-vector-based Climate model mini-app to FPGAs	Alghamdi, Moteb and Riley, Graham and Ashworth, Mike	PDPTA'21-The 27th Int'l Conference on Parallel and Distributed Processing Techniques and Applications 2021	Yes

9.4 Intellectual property rights resulting from this deliverable

N.A.

Appendices

Appendix 1 Agenda of the workshop

ESiWACE2 Second Virtual Workshop on

Emerging Technologies for Weather and Climate Modelling

October 7, 2022

10.15 - 18.45 CEST

Agenda

10:00-10:15	VC Available for test	
10:15-10:30	Welcome	Giovanni Aloisio
10:30-12:30	Session 1 – European Exascale hardware	Chair: Erwan Raffin
10:30-10:50	Leonardo supercomputing system and the national Data Valley action	Sanzio Bassini (CINECA)
10:50-11:10	LUMI - The pre-exascale system in the North	Jenni Kontkanen (CSC)
11:10-11:30	MareNostrum5 and its Data Center	Sergi Girona (BSC)
11:30-11:50	Break	
11:50-12:10	EUPEX (European Pilot for Exascale) project on the road to Exascale	Etienne Walter (Atos)
12:10-12:30	EuPILOT	Carlos Puchol (BSC)
12:30-14:00	Lunch Break	
14:00-16:00	Session 2 – Programming models and hardware interplay	Chair: Andrew Porter

14:00-14:20	PSyclone for LFRic	Iva Kavcic (Met Office)
14:20-14:40	ESiWACE2 DSLs for ICON and NEMO	Carlos Osuna (MeteoSwiss)
14:40-15:00	Vulcan FMS+ (Cancelled)	Oli Fuhrer (ETH)
15:00-15:20	Break	
15:20-15:40	Experiences with Kokkos in E3SM	Luca Bertagna (Sandia National Lab)
15:40-16:00	DestinE: opportunities & challenges for digital twins of the Earth System	Balthasar Reuter (ECMWF)
16:00-16:30	Coffee Break	
16:30-18:30	Session 3 – Machine Learning	Chair: Peter Dueben
16:30-16:50	Atmospheric Physics-Guided Machine Learning for Climate Modeling and Weather Forecasting	Tom Beucler (EPSL)
16:50-17:10	A Generative Deep Learning Approach to Stochastic Downscaling of Precipitation Forecasts	Lucy Harris (Oxford)
17:10-17:30	Building Digital Twins of the Earth for NVIDIA's Earth-2 Initiative	Karthik Kashinath (NVIDIA) and Mike Pritchard (NVIDIA)
17:30-17:50	Break	
17:50-18:10	Deep learning and differentiable simulations	Stephan Hoyer (Google)
18:10-18:30	Machine learning for weather forecasting: successes, challenges, and the future	Jonathan Weyn (Microsoft)
18:30-18:45	Wrap up and closing session	

Program Committee

Giovanni Aloisio (CMCC), Italo Epicoco (CMCC), Peter Dueben (ECMWF),
Rupert Ford (STFC), Erwan Raffin (ATOS), Donatello Elia (CMCC)

This event is funded by ESiWACE2: the ESiWACE2 project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988 - https://www.esiwace.eu/