Earth Virtualization Engines (EVE)

A draft concept paper for public comment

Conveners of the Berlin Summit for EVE*

June 5, 2023

---

*Bjorn Stevens, Natacha Bernier, Andreas F. Prein, Johanna Baehr, Peter Bauer, Hendryk Bockelmann, Sandrine Bony, David Bresch, Gilbert Brunet, Francisco J. Doblas-Reyes, Charles Ewen, David Farrell, Nicolas Gruber, Wilco Hazeleger, Torsten Hoefler, Daniela Jacob, Christian Jakob, Raghavan Krishnan, Chao Li, Jürg Luterbacher, Markku Kulmala, Pekka Manninen, Jochem Marotzke, Tim Palmer, V. Ramaswamy, Florian Rauser, Masaki Satoh, Thomas Schulthess, Jörg Schulz, Detlef Stammer, Jagadish Shukla, Julia Slingo, Adam Sobel, Thomas Stocker, Georg Teutsch, Adrian Tompkins, Tianjun Zhou

Preface

This document is being shared for community input in advance of the Berlin Summit for EVE, to be held from 3–7 July 2023. It was developed through the course of weekly virtual meetings of the core author team, starting in January 2023. The road here began for many of us with the World Modelling Summit for Climate Prediction, hosted by ECMWF fifteen years ago [1]. Ideas first articulated there subsequently developed and were expanded in what became the Flagship proposal for Extreme Earth, and later Destination Earth (or DestinE). EVE builds on the recommendation of the Scientific Advisory Board of the World Meteorological Organization to form an 'International Climate Science and Service Centre', and the Royal Society Brief on *Next Generation Climate Models* prepared for COP26. In bringing specificity to our thinking we draw on input from many recent workshops organized under the auspices of the WCRP, most directly the virtual workshop on the Future of Climate Modelling held in March (2022), a workshop on Ultra-high-resolution Modelling held in Boulder (2022), and from the development of the IS-ENES ”Infrastructure Strategy for Earth-system Modelling for 2024-2033” document. We also organized our own workshops to gain input on particular questions, notably the Lorentz Center Workshop on Digital Twin Earth, held in Leiden, NL (13–16 Feb 2023), the Yamanakako Workshop on global storm resolving models held in Yamanaka JP (12–13, May 2023), the Global Earth Observations Workshop held in Hyytiälä, FI (8–10 May 2023), and meetings organized by the Institute for Atmospheric Physics in Beijing. Our vision continues to be influenced by the emergence of a new generation modelling capability, as illustrated by projects in Japan, national projects such as WarmWorld (DE) and EXCLAIM (CH), the Horizon2020 Project NextGEMS, as well as modelling initiatives across several US agencies.

In writing this document our goal has been to capture the imagination of the climate science community to engage them in thinking about how we can work better together to advance science in the service of society during a time of profound change. As such this document is intended to provide a basis for identifying key issues and points of tension to be discussed at the Berlin Summit for EVE. A revised document will be prepared based on this input and shared in advance of the WCRP Open Science Conference in Kigali, Rwanda, in October 2023. After this final round of input a final document will be prepared for presentation at the COP28 in Dubai UAE at the end of this year.

We wish to emphasize that EVE is still a developing idea, drawing on ever-expanding yet ever-limited input. In our attempt to imagine and motivate this idea many contributions will have been left out, citations will be incomplete, and important initiatives and programmes will have been insufficiently credited or even omitted. More important than the un-honoured authority, or priority, claims is that many ideas that could strengthen our ambition may remain untapped. Hence we ask you, readers, to approach the document in the constructive spirit in which it is offered. If after reading the document you are motivated and able to join us in Berlin, please let us know how you wish to help, and we will do our best to accommodate your participation.

1 The changing need for climate information

1.1 The emerging picture of global climate change

Never before in paleoclimate records going back millions of years has our planet warmed so rapidly as today. The short and long-term consequences of this warming on humanity and the Earth-system as a whole are still largely unknown. A deeply-concerning manifestation of this change is the world's recent and repeated exposure to unprecedented weather events that has resulted in increasingly frequent, tragic and expensive natural disasters. Record-breaking heatwaves, forest and bush fires, riverine and coastal floods, droughts and sea level rise, to name only a few, are now threatening our present and future. Communities worldwide expect timely and accurate guidance to inform and support the response to imminent events and mitigate and adapt to expected changes in high-impact weather conditions in the decades and century to come. Everyone wants to know what is to be expected.

Since the 19th century, the scientific community has worked at understanding and forecasting weather, marine, and more recently climatic conditions. In the latter half of the last century, a growing understanding of the forces influencing weather and climate was complemented by the ability to observe and compute their effects, resulting in successively more physically consistent and more skillful model predictions of weather and more insightful projections of global climate change. By combining theory, observations and models, scientists have been able to document unexpected and profound changes in the global climate, and to show that we, humans, are responsible for these changes [2]. This progress was made possible by an unprecedented international effort[1] to coordinate, support and disseminate scientific work in different sectors. Community modelling efforts, such as CMIP [3, 4] and CORDEX, have been instrumental in the development and evaluation of scenarios of future change, and for modelling their impacts on a wide variety of scales. Programmes like the Global Framework for Climate Services [5] are exemplary of how to better integrate local knowledge and strengthen expertise in vulnerable sectors to build resilience to climate change. Ambitious initiatives like Copernicus Climate Services have been launched and help support the assimilation and dissemination of data from a growing observational network, notably satellites, and environmental research infrastructures. Together, such activities are greatly increasing the global human capacity to receive, critically process, and evaluate climate change information.

Thanks to these and other efforts, it is now well understood and has become widely accepted, that global warming is caused by human activities. This consensus is giving rise to policies and practices designed to limit these harmful activities, for instance by curbing greenhouse gas emissions. However, as we begin to develop programmes to confront the reality of global warming, we increasingly realize that we lack a specific understanding of what the future might hold, particularly in the Global South where long data-records are sparse, the capacity to adapt is more limited and exposure to weather and climate hazards is high. We also realize that often extreme events and their impacts are severely underestimated in their extent, or unexpected in their location of occurrence and thus, in spite of our very advanced climate science and modelling capabilities, unpleasant surprises are frequent and disruptive.

1.2 The knowledge deficit hindering adaptation and risk assessment

As changes in high-impact weather become tangible, and speculation about the potential for more catastrophic changes becomes widespread, demands for continuously updated and improved

---

[1] Notably, by international institutions such as the WCRP, WMO, UNEP and IPCC

information are mounting. Just last year, COP27[2] highlighted:

> "the need to enhance coordination of the activities by the systematic observation community and the ability to provide useful and actionable climate information for mitigation, adaptation and early warning systems, as well as information to enable understanding of adaptation limits and of attribution of extreme events".

This need is easier stated than met. While in some cases we have a general idea of what factors influence extremes, quantifying changes in ways that can help guide action has proven challenging [6]. Limitations in our ability to project how extremes will change with warming are partly intrinsic to a chaotic system, but a reliance on models designed for other purposes, or built with past limitations in computational capacity in mind, is a more immediate and surmountable obstacle. For instance, climate models that simulate decadal changes (whether that be regionally or globally) necessarily exclude the influence of processes most directly associated with many extremes, including convective storms, tidal surges, rogue waves, dust storms, riverine floods, avalanches and ice sheet disintegration. An inability to represent these and related processes means that users of such models must often contend with large systematic biases, and drifts on regional and local scales. These can be orders of magnitude larger than the changes being addressed [7, 8], and together with an absence of information on the scale of impacts, make the models ill-suited for assessing many forms of risk or for guiding action. Because adaptation efforts are costly, the remaining choice, which would be to prepare for everything, is simply not an option for all but a few.

Confronted with this knowledge gap, some countries are beginning to run and compare global climate scenarios drawing from scores of traditional global climate models, to produce legions of simulations to better identify signals of change. Others are working to develop a new generation of computationally intensive models, in the expectation, based on NWP and km-scale regional climate models, that their enhanced physical content can beat back systematic biases and better represent the local granularity of global weather and climate. Yet others, including a growing number of privately funded initiatives, are exploring the capabilities of AI (more specifically ML) to better compensate for biases of existing models, or even provide purely empirically based assessments of future hazards. These activities, as laudable as each may be on its own terms, are mostly distributed as blue-sky research efforts across a large number of research laboratories, or private companies. They are accompanied by high levels of redundancy, leading to siloed workflows and a lack of standardization, and most crucially, spreads precious talent thinly. The problem is simply too important, too urgent, and too challenging not to be approached together.

1.3 The impetus for EVE

The sense that we are better together, is the underlying impetus to EVE. In the following, we build on this impetus to identify the information needs arising from the imperative to sustainably and equitably manage our planet, and the opportunities available to address these needs. We end with a vision, or blueprint for EVE, as a coordinated global effort to tap into the potential of technology and human ingenuity to best meet the challenges identified and those yet to come.

2 New needs and a new urgency for climate information

EVE's goal is to facilitate the design and implementation of frontier climate information systems which society needs to answer critical questions. Those systems will require major advances in at

---

[2] Decision CP.27, Sharm el-Sheikh Implementation Plan

least three areas: i) simulating fundamental phenomena in the Earth-system which are are key for decision-making but inadequately represented, or even missing, from current/past generations of climate models, ii) fusing global models and observations with local granularity, using data assimilation systems designed to represent the global to local state of the climate, and iii) bringing interactivity to the application of impact models to quantitative projections of climate change.

EVE is needed to inform both mitigation and adaptation. Here *mitigation* refers to actions taken to limit climate change. *Adaptation* refers to actions taken to reduce vulnerability to climate-change-related natural hazards, to shifting weather and climate patterns, and changes in seasonality. Mitigation usually implies collective, even global, measures enacted in the form of public policy, for instance by regulating energy production, solar-radiation management or other forms of geoengineering. Adaptation often involves more local measures, for instance by greening cities to reduce the local effect of global warming, building local sea defences to manage rising sea levels, or by diversifying investments to reduce correlated risks. Additional and improved climate information is most lacking at regional and local scales, where adaptation efforts are primarily undertaken. This should not obscure important information needs related to mitigation – for instance to design new energy systems, assess high-end risks for physical and biological subsystems, or more generally inform public opinion. In addition to these ongoing traditional needs for climate information, new needs associated with managing claims of loss and damage, informing development aid, assessing capital risk, and measuring compliance with climate agreements, are emerging. These and other use-cases for climate information are discussed in more detail below.

2.1 Protecting lives and livelihoods

Every day, infrastructural or planning decisions are being made at local, regional, or national levels, in ways that risk unnecessarily exposing people and property to climate-related damages, wasting limited resources on unnecessary or even harmful efforts. Minimizing these risks requires understanding how climate will, or could, change, on the scale of the infrastructures being contemplated. Energy and water resources are vulnerable to changes in precipitation locally, on time scales that vary from minutes to decades. Ensuring food security becomes easier, or more economical and sustainable, given an understanding of changes on the watershed level in both the net amount and the temporal distribution of precipitation. Shifting rain from the growing to the fallow season has large implications for agriculture, and changes in the intensity distribution have ramifications for flooding, urban water management, and crop pest proliferation.

Those tasked with making adaptation decisions must first understand what they are adapting to. They also require an ability to assess the impact their proposed measures will have on the lives and livelihoods they aim to protect. In many, if not most cases, climate information is inadequate for informing decisions. Even when such information is available, its utility is often limited by an inability to work with it interactively to comprehend complex and indirect ramifications, which makes it difficult to explore the overall efficacy and cost-effectiveness of interventions. The difficulty of establishing standards, and the absence of procedures for certifying a greater provision of information by the private sector, hinders action and increases vulnerability. As companies are called on to disclose their risk to climate-related damages, how can due diligence in such reporting be certified?

Our degree of unpreparedness is highlighted by the recent initiative of the WMO which reports[3] that

> “one third of the world’s people, mainly in least developed countries and small island developing states, are still not covered by early warning systems.”

---

[3] Early Warning and Early Action

If, despite tremendous progress in numerical weather prediction [9], such a large fraction of the world remains so poorly informed about the threats as they emerge on the immediate horizon, how well prepared and informed are they about what climate change portends? More directly, and as the WMO points out, designing effective early warming requires climate information, which depending on the infrastructure in question, may require information on decadal, even centennial time scales. Without EVE, early warning activities will be unable to address their mandate to also account for climate change.

2.2 Measuring and monitoring

Climate change poses new challenges for measuring and monitoring. On one hand, and as alluded to below, assessing the quality and trust to be given to future projections will depend on the ability of the same systems to explain the past. Likewise, as countries and companies enter into contractual agreements to manage their interactions with the Earth-system, processes will be required to assess their compliance. A particular focus is the need to manage compliance with promised emission reductions, by linking measurements of greenhouse gases to purported emissions. Other examples include the ability to track reforestation initiatives or other interventions in land management, to assess the feasibility of potential geoengineering efforts, to identify and follow marine extreme events, or to track pollutants through the global ocean. These introduce new modelling requirements, and put a premium on understanding and quantifying the historical and evolving state of the Earth-system and making this information available to all.

Today there is an ever more urgent need to consistently integrate data from the past through the periodic reanalysis of (mostly) meteorological data. These reanalyses use physical models to assimilate observations in a way that allows one to reconstruct a physically consistent and complete description of the atmosphere (and, increasingly, surface and subsurface land and ocean) on the basis of data collected in the past. Reanalyses are, however, presently developed from data assimilation systems based on physical models developed to provide the best possible initial conditions for weather forecasts. The demands placed on an analysis (data assimilation) system in terms of monitoring the Earth-system and for verifying measures taken to mitigate climate change are different and more diverse, in particular as they must integrate local information in ways that present reanalysis systems cannot do. EVE can help ongoing activities to improve re-analyses to better introduce constraints from slowly evolving climatic variables, or to incorporate more local knowledge and data, or through access to better models more adept at using high-performance computing. This would be a boon for climate information.

2.3 Informing policy and public opinion

A deficit of specific information as to how climate will change, particularly in ways that are tangible to citizens, hinders the effort to develop a degree of consensus that is required to craft effective policies. Simply being able to better assess, anticipate, or even eventually explain, unprecedented changes will be crucial for building the needed consensus.

By creating tools that allow the public to better contextualize, for themselves, the nature of changes, and to better understand the connectivity of the world they live in, EVE offers a powerful means of advancing a global consciousness, and strengthening the ability of humans to manage the Anthropocene. Data interactivity is needed to allow users – educators, journalists, etc. – to explore possible futures and pasts, and to themselves assess the quality of information. An ability to compare projections of the past to measurements, or across independently developed models, helps assess the utility of information sources. Open and easy access to information facilitates its dissemination and encourages innovation in its use, further increasing its penetration into civil society. Open and easy-to-use systems also enable feedback to identify errors and thereby more rapidly improve. Virtualization leverages the intuition everyday people have for many of the quantities of interest, and thus builds on people's ability to relate changes to their experiences. Development and design of such systems will need to involve institutions from the national to local level, in particular in the Global South, where connectivity remains limited to cellular access at best.

2.4 Exploring unintended consequences

Expanding the frontiers of knowledge narrows the space for speculation. As climate changes, and Earth begins to enter uncharted territory, speculation, fear, and even desperation risk becoming a source of social instability. Improved scientific understanding of the potential risks of specific warming levels is needed, not just to buttress decision making, but also to better assess unanticipated or catastrophic changes to the Earth system [10]. Even the existential risk of warming imperiling the Earth's capacity to sustain humanity is not known.

An essential resource for scientists in their endeavor to increase knowledge is their ability to generate new empiricism. Programmes devoted to better documenting the past; from proxies for the deep past, to records of the recent past, from monitoring of the present, to new forms of experimentation with chemical, biological, and/or built systems, are ongoing and essential in addressing this need. A growing gap is forming in the use of computational technology to increase the pseudo-empiricism from simulations [11]. While the systematics of the existing class of climate models, suited to older generations of computers, are by now well sampled, new classes of simulations – which allow important processes in the ocean, ice sheets, land, and atmosphere, to be represented physically, rather than empirically – will introduce new systematics, increasing the empiricism, deepening the well of facts that science can draw from. As the challenge of fully using modern high-performance computers to perform such simulations outpaces the capabilities of almost all research labs, these powerful machines are, from the point of view of climate science, sitting idle. It is a little-known fact that today's climate projections rely on computational throughput that lags the state-of-the-art capabilities by a dozen years or more, and this gap is growing. How can it be that our most powerful technologies, ones with a long track record of moving climate science forward, are no longer being effectively used for this purpose?

2.5 EVE – a necessary tool for managing the anthropocene

The above arguments point out that climate change poses enormous challenges for society. It threatens lives and livelihoods, influences international relations, creates new forms of liability, and imperils our ability to envision a positive future. Fundamentally, these challenges arise from a lack of specific information about what global warming portends. Existing tools won't help fill this gap, they weren't designed for that. We need new tools. Through a concerted effort to harness the most creative minds, the most advanced technologies, and integrate these with a growing observational capacity, EVE can be such a tool, and help humans manage the Anthropocene to ensure a sustainable future for our species and our planet.

3 Developing a digital infrastructure for EVE

Impressive gains in information technology are making it possible to contemplate completely different approaches to the delivery of climate information. These and their implications for EVE are elaborated below.

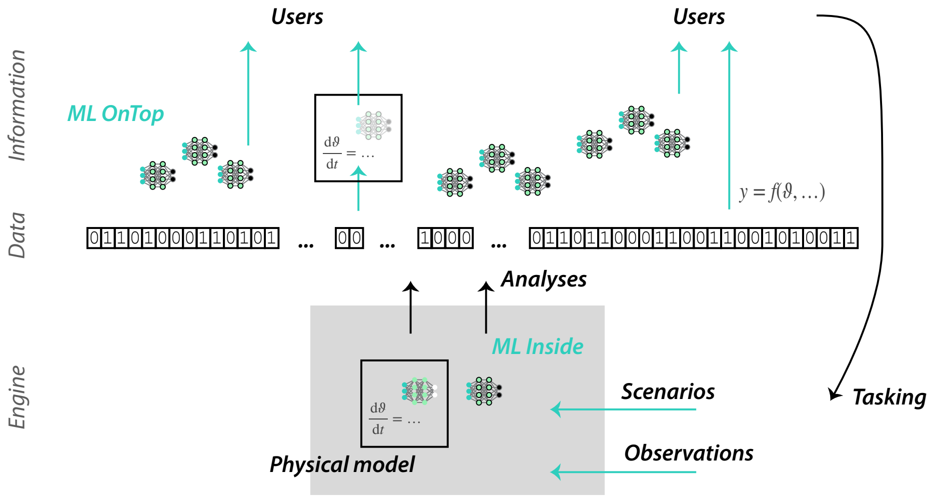

Figure 1: Schematic showing the symbiosis between ML and physical modelling, and illustrating the distinction between ML OnTop and ML Inside. Adapted from [14].

3.1 The technology landscape

3.1.1 Machine learning

Machine learning (ML), which includes deep learning among other methods, has emerged as an attractive way to create empirical models, or to calibrate otherwise poorly constrained parameters in physical models. Given sufficient training data, ML can often out-perform physical models in the space spanned by the data, especially for specific tasks. Physical models have the advantage of generalizing, which allows them to be used outside of the bounds of existing data and hence beyond the current climate state. This property allows them to be used to explore counterfactuals, such as a world without greenhouse gas forcing. The use of models as laboratories through numerical experimentation, when combined with their inherent physical 'consistency', makes them indispensable tools for creating knowledge.

Recently ML has been applied to weather forecasting, and for some quantities has already shown superior skill compared to traditional, physically-based, modelling, and at greatly reduced cost [12, 13]. Because ML is trained on *Analyses*[4] of past weather, this suggests that weather trajectories are well spanned by the evolution of previously observed weather, for new forecasts to remain in-sample. This, however, does not render physical models and observation (data) obsolete. They remain essential to provide high-quality training and initial data, as this is what ultimately will determine the skill of ML-based forecasts [14].

In a nuanced way, a similar symbiosis between physical models and ML is becoming apparent for questions related to climate. In this case, because we cannot assimilate future observations, or counterfactual climates, physical models are required to project how the system evolves under given scenarios (in the form of time-evolving boundary conditions, or for different parameter settings or model formulations). This then provides the training data that ML needs to help users interactively explore the space of simulated climates.

---

[4] The word *Analyses* is used in a more general sense than might be familiar to those who associate it with the 'analysis' of observations, for weather forecasting. Our usage is clarified in the glossary and italicized as a reminder of this distinction throughout.

For both weather and climate, the emulation of specific sub-processes whose behavior can be well constrained a priori, can help accelerate physical models [15]. In a similar vein ML can be used, and has long been used, to help assimilate observations. Although less explored, ML is also attractive as a way to compress the enormous information content produced by physical models. The more familiar use of ML, to aid in the assimilation of observations, or help more efficiently project climate trajectories following a certain scenario, we call ML-Inside. It stands in contrast to the use of ML to explore the information content of the resultant *Analyses* or projections, something we call ML-OnTop.

Distinguishing between ML-OnTop and ML-Inside helps articulate the separation of concerns between the provision of training data and the use of the data to create information and knowledge. ML-OnTop scales across what is envisioned to be an effectively unbounded application sector. ML-Inside, must be harmonized with the physical models used to assimilate observations, project scenarios, or structure output, and doesn't scale beyond a few expert communities. ML-OnTop puts a premium on the quality of the physical models, the hand that feeds them, and ML-Inside helps provide this quality as efficiently as possible. Most importantly, ML-OnTop resolves Borges paradox[5], to enable interactivity with what would otherwise be overwhelming amounts of information. In addition to the user "pull" described in the previous section, ML-OnTop is an additional structuring element of EVE.

3.1.2 High-performance computing

High-performance Computing (HPC) recently crossed a threshold that now enables global km-scale climate simulations – long used to enable more fundamental process-based studies of weather systems, and ocean circulation regionally or separately – to be used to study Earth's climate globally [16].

Models exploiting this scale of HPC are being prototyped in a few labs across the world. Performance benchmarks (Table 1) demonstrate that at their strong-scaling limit[6] a simulation throughput of 1 SYPD – a benchmark for climate studies – can be achieved using 1 PF (HPCG) are possible, on near 1 km meshes – what we refer to as km-scale. The world's most powerful computer, Fugaku, on real applications, can sustain 16 PF (HPCG) while using just under 30 MW of power. As an example, a facility that dedicated this capacity to climate studies could produce small (8 member) ensembles of 30-year ICON simulations on a 2.5 km grid, on a monthly basis using just half of the machine, matching the throughput commonly associated with much coarser resolution models.

Today, growth in computational throughput is mainly limited by limits on power consumption and by memory bandwidth, which grows slowly. While practical (HPCG) computing capacity can be expected to continue to grow, it is hard to imagine global simulations being able to access more than a factor of four increase as compared to Fugaku in the coming decade. On the model algorithmic side, dedicated engineering efforts show the potential for similar performance gains in the above-benchmarked codes. Taken together, a tenfold increase in throughput from a single facility is an ambitious, but not unrealistic, expectation for the coming decade. A facility with this capacity (160 PF(HCPG)) could generate 6000 simulated years at a 1.5 km horizontal grid-spacing, per year of operation, or closer to 30 000 simulated years at a 2.5 km horizontal grid-spacing, using just half of its computational cycles. This would enable a formidable exploration of the space of possible future climates.

---

[5] In 'On Exactitude in Science' (1946) Jorge Luis Borges explores the paradox of a map so detailed that it occupies more space than it describes.

[6] Which means that distributing the computation over a larger computer will no longer accelerate its time to solution.

Model	Δx	N_z	SYPD	PF (HPCG)
ICON	2.50 km	90/75	1	1
NICAM	3.50 km	78	1	4
NICAM	1.74 km	78	0.25	8
NICAM	0.22 km	78	0.0008	8

Table 1: Throughput from state-of-the-art global models as measured against the high-performance conjugate gradient (HPCG) benchmarks (which better account for memory bandwidth limitations). N_z denotes the number of levels in the atmosphere (and ocean when applicable).

Even finer mesh computations will still be possible but, because present simulations operate at the limit of strong-scaling, this will come at the cost of reduced throughput. As a practical matter this means that regionalization repeated widely to cover all nations (which, unfortunately, is a common national perspective) no longer confers any advantage in accelerating the time to solution. Even assuming that only some regions are to be downscaled, while it saves power, and enables larger ensembles of simulations to sample natural variability, decoupling the large scales from the evolution of the fine scales, comes at the cost of considerable information loss where it is most needed [17, 18], and ignores the fact that regional models cannot eliminate the gross large-scale biases in the global driving model.

3.1.3 The data layer

Exploiting HPC gains to provide training data, and designing systems that can make this data usable for ML-OnTop, will require substantial innovation in managing workflows and data. To appreciate the scope of the challenge, consider that a 4-byte (10 variable) state-vector of a 1 km model with 200 levels (through the atmosphere, land and ocean) is about 0.004 PB. Saving it every fifteen minutes for a year results in 150 PB of storage, equivalent to the total storage of the lustre file systems at a leading climate or weather center. Even if the data could be stored, it would be too vast to be used, creating a second version of Borges' paradox.

Resolving this paradox will require greatly reduced representations of the data, for instance using standard techniques like careful selection and compression, and novel methods including ML-inside, regeneration, and on-the-fly learning. Even then data volumes on the order of an exabyte are to be expected, and while it is one matter to store the data, it is another to do so in a way that makes it easy to access and use [19].

In addition, to minimize the expense of moving data, co-proximate compute and storage facilities will be necessary, albeit in ways that enable efficient operation across a handful of such facilities. These needs, many of which are at the scientific frontier, emphasize that EVE must have a strong component of research and development in the fields of computer science, applied mathematics, and computer engineering.

3.2 Information needs of users

The demand for km-scale simulations is rooted in the global need for local information, at greater fidelity, both to advance scientific understanding as well as to link to impacts and better integrate local knowledge, including observations. Hence EVE must serve an incredible spectrum of users, from neophytes exploring its information content, to experienced application users, to technical and scientific experts.

To mitigate against the dangers from a full-throttled implementation of HPC and km-scale simulations interpreted by ML-OnTop EVE needs to develop mechanisms to support critical inquiry and public access. Becoming overly reliant on ML-OnTop, or becoming misled by the apparent realism of simulations that too narrowly represent uncertainty, is a real danger however. Transparency in data provision is the best guard against this, and emphasises the need for EVE to provide open access to *usable* data and models. Emphasizing usability recognizes that the accessibility of data and models is much less of a bottleneck than the ability to use the data or implement the models that operate on them.

By opening access to training data (from simulation and assimilation) and building on the expertise and experience the community has gained through the Copernicus, CMIP and other similar global programmes, EVE would create a platform for innovation in ways that can more fully engage external entities (e.g., from the private sector, universities, agencies, non-profits) with EVE, to scale the provision of climate information. However, opening access to tools that only a few are able to effectively use risks increasing the inequities that open data was intended to resolve. For this reason, and to retain responsiveness to users' needs, it will be essential that those responsible for creating the training data remain in close contact with those using the data, i.e., that the small number of data providers vertically integrate across the different layers of concerns in Fig. 1, and that a degree of public access is given to both EVE's computing and storage prowess. Realizing this vertical integration will require EVE to include, as a substantial and empowered component of its staff, domain experts and boundary workers at the critical interface between EVE scientists and external user communities. Training and integration programs will be required to ensure that these experts and interfacial workers are truly representative of the global community and thus ensure that the tools meet the needs of all regions in the north and south. These boundary workers will bring local and specialized expertise to guide the selection, provision, and interpretation of the climate data, including the implementation of ML-OnTop. This will necessitate capacity building – with the positive side effects of accelerating the development of digital competencies across the breadth of humanity, and endowing the effort with global legitimacy.

3.3 EVE – the technology case

EVE's technical ambition, its data challenge, and the requirement that it be vertically integrated to support public access will require co-proximate compute and storage at large and dedicated facilities. The need to integrate local knowledge, sample uncertainty through complementary efforts, maintain access to a truly global talent pool, and establish global legitimacy, can likely only be met through regional or super-regional facilities. This suggests that EVE is best realized through up to a handful of Tier-0 facilities, distributed globally, each with co-proximate compute and storage facilities, co-located with hyperscaling data centers. This presents challenges for system design. For instance, methods must be developed to ensure that models are portable across infrastructures, including generic (or cloud) compute resources. Likewise, training data must be standardized and globally accessible. Implementing co-proximate methods that work across the federation, and protect intellectual property implemented at and across even a small federation, raises additional challenges that EVE must surmount.

This points out how EVE, as a grand technical challenge, is also a tremendous opportunity. In creating virtual Earths EVE will spur technological innovation, its data provision will nurture new economies, and its use will strengthen digital competencies.

4 Advancing the science to support adaptation and mitigation

As climate change unfolds, climate science is struggling to keep up with societies' growing need for information. Modelling has, and continues to be, a motive force for science, with progress most limited by the inability of models to resolve key Earth-system processes. However, many of these processes can now be resolved in km-scale climate models offering breakthroughs in major and long-standing scientific challenges. Examples include quantifying the changing frequency and intensity of extreme events; the effects of storm clustering on Earth's radiant energy budget [20, 21], how precipitation is distributed in space, time, and intensity; the impacts from representing the full spectrum of precipitation as they reverberate through the Earth-system [22, 23]; the nature of interactions between Earth-system components [24–26]; and the likelihood of tipping points. This gives urgency to the need to fully develop and apply such models to the study of climate change [27].

4.1 The scientific case for km-scale Earth-system models

Higher resolution is not an elixir, but when it comes to improving Earth-system modelling, nothing comes closer. This has been evident in the CMIP process, whereas a subset of models have slowly edged toward scales of tens of kilometers, their systematic biases have been reduced [28–31] and important new sources of predictive skill have been identified [26, 32]. There is mounting evidence that further substantial improvements will arise from the systematic application of yet finer, km-scale, models. Kilometer-scale models have shown a capability to better represent the global tropics, from eliminating the double peak in tropical convection [5, 33], to improving the simulation of tropical wind patterns [34], to capturing major modes of tropical variability that conventional models miss [33, 35]. In the extra-tropics fine scales improve the representation of the storm tracks [36], better represent orographic drag [37], and reduce precipitation biases [38]. Across Earth, km-scale models improve the representation of precipitation amount, intensity, duration, and phase [39], consistent with their better resolution of precipitation extremes [40, 41], and are beginning to identify responses to warming not seen in lower resolution models [42].

Discussions of resolution often focus on the atmosphere, because the scale and speed of processes operating there make it computationally demanding. This means that it is the atmosphere that limits the application of higher resolution to other components of the climate system. This focus sometimes obscures what practitioners have long known, which is how important high resolution also is for the other components of the Earth-system, both individually and as part of the coupled system. Kilometer-scale simulations have been shown to resolve the erroneous simulation of freshwater storage in snow [43, 44] and ice [45], which is important for hydrology. They improve land-atmosphere coupling [24, 25], and the impact of sub-surface hydrology on extremes [46]. Ocean-atmosphere coupling is likewise sensitive to km-scale processes [26], and even in the absence of an atmosphere, meso-scale eddies captured by km-scale ocean models strongly influence how the southern ocean responds to changes in wind-stress [47]. Along with smaller eddies and fronts, they have been highlighted as being ”key” for how the ocean and cryosphere evolve under climate change [18, 48]. Links between the land biosphere and ocean biogeochemistry are influenced by riverine elemental inputs, shelf carbon dynamics, and sub-mesoscale circulation whose transient dynamics require km-scale resolution to be represented [49]. The stability of ice sheets has been shown to depend on km-scale features in its grounding line [50], and in the representation of the surrounding ocean [51].

Because of computational limitations, km-scale models have mostly been applied to regional domains, or to global grids for short time periods, and with limited coupling of Earth-system components [52]. Leveraging the full advantages of km-scale models will require these efforts to

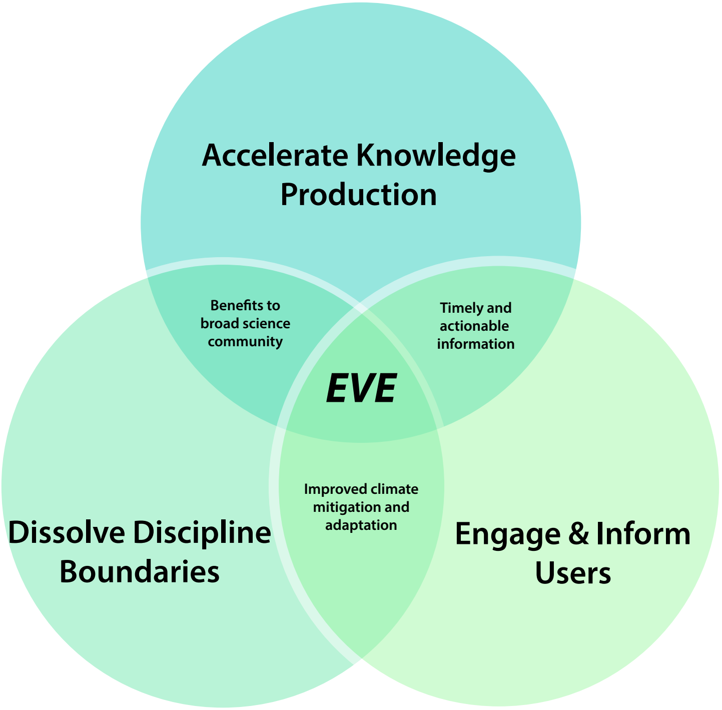

Figure 2: Successfully realizing EVE will 1) accelerate scientific discovery and digital innovation; 2) lead to new discoveries by dissolving boundaries across research disciplines; and 3) revolutionize our ability to inform and engage with users on a global scale.

be accelerated and extended. This is generating considerable excitement, or anticipation, as since the dawn of climate modeling, each addition of a new class of numerical model to the existing hierarchical suite has produced a step change in our ability to understand and simulate the climate system. Today, by representing the full spectrum of interactions from the convective scale to the mesoscale to the planetary scale, across all components of the climate system, km-scale models promise needed breakthroughs in climate science, and can reshape the field in three important ways (Fig. 2).

1. **Accelerate knowledge production**: Kilometer scale models will introduce new empiricism, and facts. These will stimulate scientific activity. The new paradigms that EVE must develop to run the models, distribute and analyze their data will also greatly advance digital competency.

2. **Dissolve boundaries**: By working on the same spatial scales as our sensor footprints, and as studied by national meteorological, hydrological, and environmental agencies EVE will dissolve long-standing boundaries between these fields. New discoveries will emerge from the ability to explore the coupling between Earth-system components (riverine carbon transports and ocean biogeochemistry, land-surface and sub-surface interactions, ice-sheet interactions with grounding lines, etc.) globally.

3. **Engage users**: By describing Earth's climate on the scale of transient dynamics that people experience, EVE will open new opportunities to engage and inform existing users, and at the same time begin to engage a much broader user community, generating new knowledge economies [14].

4.2 The scientific challenge of km-scale

To capitalize on EVE's potential will require solving not only large technical, but also scientific, challenges. While EVE can make initial progress by leveraging and further developing existing modeling systems, it will need to nurture and sustain an ongoing programme of model development tailored to the scale of the problems and the breadth of the user communities that it is targeting.

4.2.1 Model building

The necessity of developing km-scale models should also not hide the scientific challenges that this poses. Experience in the development of regional (as opposed to global) km-scale models suggests that, despite immediate and significant progress in simulating many aspects and phenomena of the Earth-system, challenges will inevitably arise in simulating other aspects [53, 54], an area, where if carefully formulated in ways that generalize, is proving amenable to the combination of bursts of finer scale simulations and ML [55, 56]. Improvements in many model components, (e.g., land surface processes and hydrology) will require expertise from communities previously only marginally engaged in Earth-system modelling. Model deficiencies due to imperfect representations of sub-km scale physical and chemical processes (e.g., cloud processes, atmospheric and oceanic turbulence, ice-fracture), and from representations of biological processes (marine ecosystems, and the terrestrial biosphere) at all scales, will need to be addressed with the appropriate level of precision. Solving these problems, by bringing expertise in modelling, information theory, observations and theory together to work on societally relevant scales, will accelerate knowledge production and transferability.

4.2.2 Uncertainty estimation

EVE offers new opportunities to better sample uncertainty. Model diversity is presently understood in terms of differences that arise from different implementations of essentially the same model.[7] By creating a new class of models, EVE makes a major contribution to diversifying our epistemological basis. Moreover, as a base-funded and coordinated activity, EVE can strategically sample how its own implementation strategies affect the systematics of its model ensemble. At present, this must be addressed in a much more ad hoc fashion, which introduces large inefficiencies.

It is sometimes feared that the computational intensity of km-scale models will limit the ability to sample scenario uncertainty, or natural variability. To the extent existing models sample this well, km-scale models will simply correct the large biases in their low-order statistics. If existing (CMIP) models do not realistically sample the higher-order statistics of the climate system, EVE will introduce a new way to do so. At the same time it introduces the needed training data (empiricism) to make ML-based ensemble inflation techniques [57] more effective, and the technical foundation for exploring more novel methodologies.

4.3 EVE – the science case

Like the Apollo project, EVE must harness tremendous technical and scientific capability to provide a specific service, and in doing so, help us see the world in new ways.

---

[7] Traditional climate models are all based on fundamental and common assumptions, namely that the intermediate scales play no role in climate, that the small scales couple directly to the large scales. While these assumptions are known to be incorrect, the premise is that deficiencies that arise can be overcome using different ways of compensating for the errors, which is the basis for large multi-model ensembles.

To be successful EVE must maintain and support a focused research and development programme, both to improve and implement models and to understand how and why they behave as they do. These must have strong connections to external efforts, perhaps through the sponsorship of blue-sky research on selected topics at external laboratories. In addition to a strong scientific focus, EVE must support basic advances in technical fields, such as informatics, numerical mathematics, and especially AI. Doing so will ensure that EVE stimulates science, and knowledge gains to maintain pace with our changing world.

While primarily intended as a telescope looking into the future of our own planet, EVE can also help us see ourselves through the lens of other planets. For instance, as on Earth mesoscale processes appear to also play a fundamental role in planetary atmospheres (e.g., Mars [58], Venus [59], Saturn [60], Titan). Kilometer-scale modeling can thus help to prepare space missions and exploration, and to interpret new observations of the solar system and of exoplanets. To enable this the EVE science and infrastructure should be coordinated in ways that facilitate its use by the planetary science community.

5 Optimal Earth-system data integration and monitoring

Coordinated and open Earth observations enable decision-makers around the world to better understand the issues they face, and better shape policies. Together with simulation, observational data is fundamental for advancing Earth-system science.

A feature that distinguishes observational data from model output is its heterogeneity. For Earth observations, intrinsic heterogeneity, across dimensions like type, quality level, redundancy, and sampling, is convolved with the space-time heterogeneity of the Earth-system itself. It is what makes standardizing access and enabling interoperability of observations and software especially challenging. Initiatives like the FAIR Data Principles [61], address some of these issues, but in so far as they apply to the use of data, they overlook a much larger problem. Most data is not used.

Insufficient, or inadequate, use of data is partly a distribution problem, but even in cases where, after a tremendous effort, data is made readily accessible, its effective use is far from guaranteed. Of the more than 100 million daily observations coming into most advanced operational NWP centers, less than 1/3 is assimilated. Reasons for the non-use of data are varied, and can be due to economies produced from dealing with instrument and sampling uncertainties, challenges of maximizing information content of quantities of interest, as well as major difficulties in interpreting complicated observations. Our present state-of-the-art assimilation approaches are built around a linearized, large-scale, balanced view of the atmosphere and are ill-suited to ingest nonlinear, high density, remotely sensed measurements at the km-scale. EVE will need to support the science to move effectively towards such km-scale assimilation systems that can leverage the observation network to its full potential. At the source of many of these challenges are limitations in the realism of the models used within the data assimilation system. As scientific instruments are becoming more sophisticated and measurement devices become easier to deploy – from commodity devices in cities, farms, cars and mobile phones and increases in connectivity aided by *in situ* learning, e.g., tinyML – the problem of non-use, and the lost opportunities that come with it, are fast-increasing (Fig 3).

5.1 Strengthening data assimilation

Regarding the fundamental problem of combining model simulations with observations within an objective framework, data assimilation has become an indispensable tool for making data usable, so much so that most people interact not with data, but with the state of a model constrained by the data it assimilates — usually in the form of reanalyses. These reanalyses are extensions of the analyses used to initialize weather forecast models as they apply the same model and assimilation configurations to long periods such that the monitored variability is the sole result of weather evolution and observation availability. Reanalyses offer the opportunity to incorporate data that might not have been available within the time-critical window during which the analyses must be performed, and can therefore recycle historical data with modern methods. These are important and minor adjustments to a painstaking and labor-intensive system that has, over decades, been optimized to improve weather forecasts. They also make the data more usable for climate studies, but highlight how the assimilation of Earth-system data is often not the first priority, particularly as no agency has the mandate to build assimilation systems that are optimized around the specific needs of Earth-system monitoring or to help evaluate and improve climate models. As we strive to make better climate predictions we have come to appreciate the importance of data assimilation as a reference methodology to bridge between models and observations. At the same time, we are beginning to appreciate the importance of advanced reanalyses, and their co-provision with data that might not be assimilated, for training deep-learning methods geared at weather prediction emulation and for modelling impacts.


EVE will help to better structure high quality and high temporally resolved one-off measurements from different sources and lead to better assimilation.

Sophisticated new satellites provide observing quantities and qualities that go beyond the current generation of weather and climate models.

Marine measurements will be important to complement information from land and space.

Ground based networks from islands stations such as those from ACTRIS infrastructure will provide high-quality data and information on short-lived atmospheric constituents and on processes from natural and controlled atmospheres.

Figure 3: Observations come in many forms. EVE provides a framework for helping data providers, and providers of data *Analyses* (including re-analyses) deliver their products in ways that ensure interoperability with software simulation, and minimize barriers to their effective and equitable use.

EVE's push to advance our understanding of the climate system globally, with local granularity, creates new opportunities to strengthen the efforts of established data providers and designers of data assimilation systems to provide improved data products, and to lower barriers to their effective and equitable use. Global models with local granularity are particularly attractive for assimilating data in complex terrain not resolved by traditional models. Combining this capability with ML and an Earth-system, rather than just meteorological, focus also widens the scope of interesting observations, (e.g., to include data from the Internet of Things, privately owned remote sensors, etc.). EVE, to the extent it encompasses an ambition to support Earth-system data assimilation (around ongoing efforts) will create new data needs, and new opportunities to make more effective use of old data.

As alluded to above, all data cannot and should not, or not only, be assimilated. EVE can advance the use of data, by helping provide access to primary data, for assessment and verification purposes, and by ensuring the interoperability of models, data and software. Doing so will add value to observations from the Internet of Things. For example, by lowering barriers for linking measures quantities like crop losses from freezing rain, blowdowns of tree stands from storms, images of damages within the built environment, coastal erosion from storm surges, or reports of complications from asthma in association with dust storms EVE opens the door to new discoveries, and to a greater and more intensive use of data. Interoperability further aids virtualization and ML [62] to create powerful tools for leveraging observations from super-sites to design experimental studies, better optimize the observational network, or just allow citizens to virtually explore their planet's past.

5.2 Strengthening data collection and provision

5.2.1 Reinvigorating ground-based *in situ* measurement system

In situ (ground-based, airborne, waterborne, and shipborne) measurement stations require constant maintenance and costly infrastructure. As remotely sensed observation increase, this places greater pressure on optimizing what inevitably will be a reduced amount of *in situ* measurements. To meet global grand challenges (climate change; biodiversity losses; air, water and soil pollution; water, food and energy supply; acidification, deforestation, and pandemics to name only a few) which all are very interlinked, we cannot have too much data. Often this requires the types of measurements, and the coincidence of measurements, that cannot only be made from space, or which might not be immediately relevant to weather and climate forecasting. During recent years, global and European *in situ* Research Infrastructures have been developed. [63]. As an exemplary case, SMEAR (Stations Measuring Earth surface Atmosphere Relations) concept, developed since 1991, is designed to provide the capacity (data volume, number of measured parameters, connections to remotely sensed observations) to answer multidisciplinary scientific and societal questions, even those which do not exist or which we cannot anticipate today [64].

As many climate questions overlap and complement other research challenges, diverse arrays of measurements will continue to be needed, new measurement systems will need to be designed, and gaps in coverage will need to be identified and filled. For instance, a lack of open observations throughout the Global South is crucial to fill. EVE must work with agencies across regions to ensure the value of integrating existing data into EVE's open access platforms is clear, and to collaborate to ensure the tools provide demonstrable leveraging of the information for local climate applications and questions. To ensure this, national research and weather agencies should be integrated from the outset so that EVE pays more than mere lip-service to open data access in the Global South, or that these aims are not relegated to an afterthought. At the same time, it will be essential to optimize network designs, so limited resources are effectively deployed. EVE's emphasis on the interoperability of data and software, can support communities in this process of optimizing ground-based data collection.

5.2.2 Essential Climate Variables (ECVs)

It will not be possible to observe everything everywhere but, what we observe we should make usable, and what we want to observe, or perhaps think we no longer need to observe, should be decided as rigorously as possible. These are questions that the Global Climate Observing System addresses through its definition of Essential Climate Variables (ECVs). Ongoing activities dedicated to further developing and testing ECVs, will benefit from a more seamless hierarchy of stations/observations (comprehensive flagships, flux stations, low-cost sensors utilizing 5G-7G networks) to assess the availability and utility of additional variables. To take full advantage of these investments, however, it will be imperative to develop tools that can help maximize the use of the data. This includes the FAIR principles, but as alluded to above, this is only a beginning. EVE's capability to complement these measurements with km-scale local models, and even finer (dm to hm) scale simulations for local process studies can help optimize the expensive deployment of research infrastructures and their associated data curation.

Consideration also needs to be given to unfamiliar data sources, from remote areas, which may be of fundamental importance, but not immediately associated with impacts. For example, sub-surface flows around ice shelves, or deep ocean heat content may have a relatively small impact on weather prediction systems, or on economic productivity in the short term, but their absence can leave us ignorant of changes that could portend risks, or a high potential for climate-related damages. Likewise, observations should be collected based on their ability to inform error estimation. By increasing the proximity between observations and model output, and ensuring the interoperability of data and software, opportunities for data analytics, knowledge transfer, and ultimately impact assessment will also increase.

5.3 EVE – the observational perspective

EVE obviously cannot substitute for the wealth of ongoing activities dedicated to data collection and dissemination. Through improved global models with local granularity it can, however, support and strengthen ongoing efforts to improve the design of observational networks, to intelligently choose essential climate variables, and to improve the capabilities of data assimilation. Likewise, it can strengthen the use of data to improve the quality of the modelling and the ability of user communities (including scientists) to assess model-based projections. Improved system design and data selection will be helpful, improved assessment of projections is essential. To meet these goals EVE must emphasize the co-location and standardization of data and its interoperability with software, whereby the former includes observational data and model output; the latter includes analytics, learning, and simulation systems.

6 Elements of structure and governance

6.1 EVE's climate information mandate

To make a difference it will be essential that EVE is structured to continuously provide and update climate data – *Analyses* of the past and future scenarios – in ways that minimize barriers to use, and lead to the widest possible dissemination of its associated information. EVE's primary mandate must be to improve and increase the empiricism, through its improved *Analyses* scenarios (projections), in ways that support the development of accompanying storylines or narratives, to inform and motivate action. User needs should be the main driver for updating *Analyses*. Ensuring the data, information, and knowledge associated with these *Analyses* are rapidly, routinely, and robustly delivered to society, will require operational rigor, which is something that present (research-based) efforts lack.

The nature of this rigor can draw inspiration from numerical weather prediction, which produces forecasts on a schedule, and includes a regular update cycle that harmonizes new model developments with HPC capacity growth, rigorous testing, and new service provision. Updated weather forecasts and dissemination services are then routinely provided using updated production chains to address growing and changing user needs. For EVE, update cycles need to be defined by their ability to increase the empiricism of its climate data. In practice this would be measured by the value updates add to what, must effectively be seen, as training data. This will require coordination with other EVE nodes and the user community – as the trade-off between a better sampling of uncertainty versus improved information quality is challenging to assess and vexes the present provision of climate information. Critically, governance mechanisms must be developed to enable the boundary workers who interface with external users to carry those users' concerns back to influence the governance and scientific direction of EVE.

6.2 EVE as a North-South partnership

In developing regional nodes, and given the need for co-proximate compute and store facilities in proximity to hyper-scaling data centers, EVE will, by definition, provide remote access for almost all of its users. This creates opportunities for it to develop as a true partnership between the global North and South. Partners would participate in the governance of the node, co-develop and co-design tools, and have privileged access to its computing and storage resources. Access to standardised tools and data, including ML and training data [14] will also help EVE dispel inequities that arise from asymmetric capabilities to make use of open data. Most importantly, partners would help design and staff training programmes that link to ongoing activities in the region – for instance in cooperation with regional climate centers and in conjunction with the regular climate outlook fora – to support capacity building and increase bandwidth amongst the partners. These should also leverage existing partnerships and regional networks such as the pan-African university and the WASCAL pan-west African PhD program to give two Africa-focussed examples.

Taking advantage of the experience of many institutions and plethora of new innovate educational tools gained and developed during the COVID pandemic, a training programme could involve the development of online courses with regional climate centers to target a very broad community of users and disciplinary expertise. Past experience of pan-continent projects has shown that small-scale training programmes with attention focused on few individuals does not work, as it leaves training vulnerable to staff movements, and can often fail to establish self-sustaining local capacity. By employing high quality, massively parallel online training components as the first step EVE aims to achieve strength in breadth rather than depth. From enthusiastic and talented participants of the online courses, scholarships to training programmes could be offered. The top graduates of these programmes would then qualify for an in situ intensive training (with potential certification at masters level) at one of EVE's partner centers followed by a rotating (5-year) staff position at an EVE node. This would allow a considerable fraction of the people employed by EVE to be drawn from the Global South, building bandwidth among the partnering institutes and strengthening their ability to make effective use of EVE's resources in the long run.

6.3 'Where are the Humans?' – Ethical, legal, and social issues

EVE raises profound ethical, legal and social issues. Data streams that might be intended to help evaluate future projections, or which might be useful for measuring compliance, can be misused to support undesired monitoring. Asymmetries in the ability to use information may, through its free provision, inadvertently lead to an increase in inequality. The difficulty users may have in contextualizing information about the future, or the possibility that some users will manipulate this information in ways that are deliberately misleading, must be contended with. As increasingly detailed information is provided across societal sectors, who bears the liability for unintended consequences associated with the use of this information? How can intellectual property be protected across EVE systems? How can legitimacy be established in ways that allow EVE to function internationally? To address these, and many other questions, EVE must meaningfully engage and continuously involve experts from the social (behavioral to political) sciences.

6.4 Initial thoughts on governance

EVE's structure should be based on a distributed framework to preserve national or multinational sovereignty, make the best use of existing infrastructures and programmes, and respond with agility to changing economic boundary conditions. Ideally, EVE is optimised for synergistic developments without creating unmanageable dependencies. The synergy must leverage the expertise and resources of multiple partners. For this, EVE needs to establish a governance framework that can quickly spin up from existing seed activities where they exist (for example DestinE in Europe) but also be responsive to new opportunities. As sufficient funding is essential to create critical mass, EVE needs to draw from existing funding programmes but stimulate substantial new funding streams at the same time. The latter will need sufficient lead time.

Governance includes science, digital technology, funding and process management aspects, most of which have inter-dependencies. The governance framework needs to be based on co-development and cooperation principles through research centres, HPC centres, operational agencies and service providers but also private entities from different countries.

The role of public-private partnerships requires specific attention, in particular for establishing EVE's digital technology agenda including the provision of supercomputing resources. EVE is also the platform for a public-private symbiosis on AI/ML that empowers the public sector through entirely new tool sets but also enables business development in service sectors. Such a symbiosis will be instrumental for democratising software and data, and sharing and co-developing cost-effective information systems and decision-making tools for those countries that are most exposed to change and least enabled through classical prediction systems.

Governance implies a regulatory framework and multi-party agreements on roles and responsibilities, how to coordinate individual developments and transition research to production. It also needs a concerted approach to data and software licensing and management. Transparency and traceability, accessibility and sustainability of software and data are critical, and existing/accepted best practices, regulations and policies need to be ingested as much as possible. The EVE implementation plan will need a dedicated effort to assess governance options that are functional in the science-technology-policy complex and that is based on lessons learned from inter-governmental organisations like ECMWF, ITER, IRENA, and CERN, subscription-based programmes like ESA's Earth-observation programme, and public-private partnerships, for example, created for dealing with COVID-19 with framework support from the World Bank.

7 EVE – an initial blueprint

The world is warming. This warming is expected to have catastrophic consequences for segments of humanity, but for whom and how many is unknown. Efforts to mitigate against the warming are motivating investments in the trillions of Euro to restructure economies and energy systems, similar investments are being contemplated to adapt infrastructure to increase resilience, and massive and risky interventions are being discussed, with little idea of the sustainability or efficacy of such responses. Never before have global societies faced such a profound and common threat. Yet our response remains inadequate and atomized. Yes, across agencies and institutions climate is a high priority. However, with the exception of short-term funding for academics, whose impact and efforts disproportionately target small circles of policymakers, climate change remains, at best, the second most important problem for even the most engaged agencies, and (with rare exceptions) lacks a meaningful international dimension. We must do better, and we must do better together. EVE, which is envisioned as a digital infrastructure that harmonizes the latest advances in information technology (HPC and AI) to establish and maintain a global km-scale

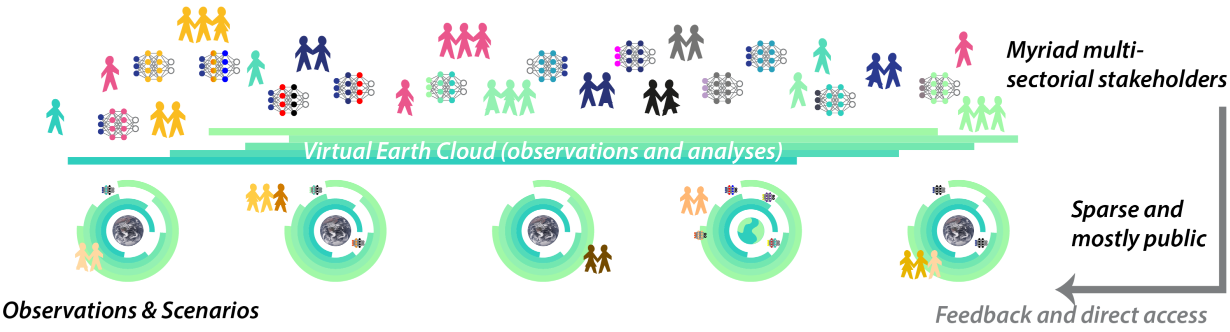

Figure 4: Schematic of EVE concept, highlighting the three layers whereby Earth Virtualization Engines (the engine in Fig. 1 create *Analyses* that, together with primary data, populate the Virtual Earth Cloud (the data layer) and make it available in ways that minimize barriers to use allowing information to penetrate society more easily.

climate prediction and information system — a Virtual Earth Cloud — across regional nodes, outlines how we can do so, and how in doing so we can help humanity positively engage in first imagining and then shaping its own future.

EVE's climate prediction and information system would aim for three levels of information provision: (i) routine (effectively operational) provision of the best possible *Analyses* of future trajectories of the Earth-system; (ii) provide and maintain a digital infrastructure, a Virtual Earth Cloud, to enable open equitable and *effective use* of these *Analyses*; (iii) service provision to maximize stakeholder engagement and public benefit. In this context *Analyses* refers to global km-scale spatial and sub-hourly (15-30 min) temporal description of the state of the Earth-system. To provide the best possible *Analyses* will require EVE's weather and climate data to be quality-assured, to include uncertainty quantification, and be continuously updated to reflect the evolving state of the science and digital infrastructure. *Effective use* of the Virtual Earth Cloud means enabling interactivity by minimizing barriers to exploring and using the cloud, including its primary data, e.g., observations and models upon which the secondary data is, or could conceivably be, based.

To address these aims EVE nodes must:

- be vertically integrated across its three levels of service provision to support user interactivity;
- operate dedicated Tier-0 compute and storage capabilities to deliver km-scale Analyses;
- engage and support under-resourced partners;
- be autonomously constituted and governed subject to minimal rules of confederation chosen to ensure a unified, and globally aligned, information system.

Vertical integration refers to the need to support technology development and research in technical fields (computer science, computation science, computer engineering, and applied mathematics); in Earth-oriented domain sciences (meteorology, oceanography, biogeochemistry, hydrology, land surface, and cryospheric processes), in application areas (agriculture, health, land and water management, urban planning, mobility, etc.), and in the boundary lands between these areas and civil society. The breadth of the required expertise cannot be maintained in a single center, but each node should aim for a critical mass that allows it to maintain its operations while engaging with experts and agencies in each of these areas.

Support for under-resourced communities must include programmes to identify, attract, train, and eventually employ staff members from these same communities and thereby engage them as active partners. As an example, an EU-initiated EVE node could be constituted as a partnership between the EU, West Africa and the Caribbean to create outreach and training programmes, satellite compute and data facilities, and rotating staff positions and thus give substance to this partnership.

Tier-0 compute and storage facilities can be defined by their power envelope, which based on current trends implies 30 MW or more. These centers already exist, EVE simply proposes that some subset of them be dedicated to managing climate change. Concentrating compute resources in a few high-power facilities makes power use more visible, but allows for more efficiency and better-coordinated access to carbon-neutral sources. The facilities should be large enough to provide *Analyses* with km-scale granularity while reserving at least half of their resources for research, development, and application activities, e.g., uncertainty quantification, model testing and calibration, training on the *Analyses*, etc.

Given different national circumstances, and the need to deliver local solutions to local problems, EVE nodes should be constituted and governed autonomously, with rules of confederation kept to a minimum while ensuring the global success of the project. These rules of confederation should principally concern themselves with: (i) maintaining the interoperability of models and data across the Virtual Earth Cloud that they create; (ii) defining a common strategy for populating the Virtual Earth Cloud with *Analyses*; and (iii) supporting the necessary exchange of information to ensure the efficient development and dissemination of adaptation solutions.

In practical terms, we envision that EVE would be well served by a federation of 3-5 regional nodes, with each node comprising a core staff of 200-300 individuals[8] with strong and symbiotic links to academia, national research laboratories and programmes, and public agencies (climate services) in partner countries. We would recommend its provision of *Analyses* to focus on temporal windows of 50 years about the present, with fewer centennial-scale windows to assess high-end risks. A funding volume, of € 300 M per year, per node could meet these targets. With 3-5 nodes this would comprise a yearly expenditure of € 0.9 B - € 1.5 B over an initial period of 7-10 years with an evaluation after five years. This amounts to a global, ca € 10 B investment to enable interactions with virtual Earths, to help us understand, manage, and as necessary adapt, to a warming world.

---

[8] These would operate the backbone of an EVE node to develop the models, maintain the infrastructure, and support the application and user, obviously this could differ drastically from node to node given the individual circumstances.

Glossary

AI Artificial Intelligence. 2, 21

Analyses This term is historically associated with the best estimate of the state of the weather based on observations, later it was generalized to include the consistency with a model. We use it generically here to refer to the best description of the state of the system given some conditions, i.e., a scenario of forcing, past information, consistency with a physical model, etc. In our usage *Analyses* can also be performed for a scenario, in which case it is synonymous with a projection. Our more general usage of the term allows us to emphasize the similarity in how we might provide data for the past, present and future. 6, 7, 14, 16, 19, 20

CERN European Organization for Nuclear Research. A research laboratory to uncover what the universe is made of and how it works. 18

ECMWF European Centre for Medium Range Weather Forecasts is both a research institute and an operational service that is funded by member states to provide global numerical weather predictions and other data. ii, 18

ESA European Space Agency. Its mission is to shape the development of Europe's space capability and ensure that investment in space continues to deliver benefits to the citizens of Europe and the world. 18

HPC High performance computing. Usually reserved for machines that are in the top forty or so (Tier-0 or Tier-1) of computing power, which presently spans about 0.2 PF(HPCG) to 29 PF(HPCG). 7–9, 16, 18

HPCG High performance conjugate gradient, a practical performance measure for the physical models used in Earth-system simulation. 7, 8, 21

ICON ICON stands for icosahedral non-hydrostatic. It refers to a coupled model of the Earth-system developed by a growing consortium of European partners, with support for components representing vegetation dynamics, cryospheric processes, the carbon cycle, and air quality that can be run at km-scale (1.25 km) resolution globally. 7, 8

IPCC Intergovernmental Panel on Climate Change. A United Nations body for assessing the science related to climate change. 1

IRENA International Renewable Energy Agency. A global intergovernmental agency for energy transformation with 137 member states plus the EU. 18

ITER International Thermonuclear Experimental Reactor. A thirty-five nation collaboration supporting the worlds largest fusion experiment. 18

ML Machine Learning – A specific form of AI, which for many applications exploits a specific technique which is called deep learning. 2, 6–9, 12, 14, 15, 17

NICAM NICAM stands for non-hydrostatic icosahedral model. It refers to an atmospheric model developed by a consortium of Japanese partners. It was the first of a new class of km-scale atmospheric models, and has been used to simulate the global atmosphere at 220 meter scales. 8

NWP Numerical Weather Prediction. 13

PB peta (10^{15}) byte: Large data centers for weather and climate typically encompass hundreds of peta-bytes of data on fast storage, and much more on tape. 8

PF peta (10^{15}) FLOP (floating point operations per second). 7, 8, 21

SYPD Simulated years per day, a measure of computational throughput. 7, 8

UNEP UN Environment Programme. 1

WCRP World Climate Research Programme. The WCRP is in turn sponsored by the WMO, the International Science Council (ISC) and the Intergovernmental Oceanographic Commission (IOC) of UNESCO). ii, 1

WMO World Meteorological Organization. 1, 3, 4, 22

References

1. Shukla, J., Palmer, T. N., *et al.* Toward a New Generation of World Climate Research and Computing Facilities. *Bulletin of the American Meteorological Society* **91,** 1407–1412. ISSN: 0003-0007, 1520-0477. (2022) (Oct. 2010).
2. IPCC. *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* (Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 2021).
3. Meehl, G. A., Boer, G. J., Covey, C., Latif, M. & Stouffer, R. J. Intercomparison Makes for a Better Climate Model. *Eos Trans. AGU* **78,** 445. ISSN: 0096-3941. (2020) (1997).
4. Eyring, V., Bony, S., *et al.* Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization. *Geosci. Model Dev.,* 1937–1958 (2016).
5. Heim, C., Leutwyler, D. & Schär, C. Application of the Pseudo-Global Warming Approach in a Kilometer-Resolution Climate Simulation of the Tropics. *Journal of Geophysical Research: Atmospheres* **128,** e2022JD037958 (2023).
6. Pitman, A., Fiedler, T., *et al.* Acute climate risks in the financial system: examining the utility of climate model projections. *Environ. Res. Climate 1* (2022).
7. Palmer, T. & Stevens, B. The Scientific Challenge of Understanding and Estimating Climate Change. *Proceedings of the National Academy of Sciences* **116,** 24390–24395. (2022) (Dec. 2019).
8. Slingo, J., Bates, P., *et al.* Ambitious Partnership Needed for Reliable Climate Prediction. *Nat. Clim. Chang.* **12,** 499–503. ISSN: 1758-6798. (2022) (June 2022).
9. Bauer, P., Thorpe, A. & Brunet, G. The Quiet Revolution of Numerical Weather Prediction. *Nature* **525,** 47–55. ISSN: 0028-0836, 1476-4687. (2020) (Sept. 2015).
10. Stirling, A. Keep It Complex. *Nature* **468,** 1029–1031. ISSN: 1476-4687. (2023) (Dec. 2010).
11. Stevens, B. & Lenschow, D. H. Observations, Experiments, and Large Eddy Simulation. *Bulletin of the American Meteorological Society* **82,** 12 (2001).
12. Lam, R., Sanchez-Gonzalez, A., *et al. GraphCast: Learning Skillful Medium-Range Global Weather Forecasting* Dec. 2022. arXiv: `2212.12794 [physics]`. (2023).
13. Pathak, J., Subramanian, S., *et al. FourCastNet: A Global Data-driven High-resolution Weather Model Using Adaptive Fourier Neural Operators* Feb. 2022. arXiv: `2202.11214 [physics]`. (2023).
14. Bauer1, P., Dueben, P., *et al.* Deep Learning and the New Economies of Weather and Climate Prediction. *Nature Reviews of Earth And Environement* **XX,** submitted (2023).
15. Meyer, D., Hogan, R. J., Dueben, P. D. & Mason, S. L. Machine Learning Emulation of 3D Cloud Radiative Effects. *Journal of Advances in Modeling Earth Systems* **14,** e2021MS002550. ISSN: 1942-2466. (2023) (2022).
16. Hohenegger, C., Korn, P., *et al.* ICON-Sapphire: Simulating the Components of the Earth System and Their Interactions at Kilometer and Subkilometer Scales. *Geosci. Model Dev.* **16,** 779–811. ISSN: 1991-9603. (2023) (Jan. 2023).
17. Wing, A. A. & Emanuel, K. A. Physical Mechanisms Controlling Self-Aggregation of Convection in Idealized Numerical Modeling Simulations: SELF-AGGREGATION MECHANISMS. *J. Adv. Model. Earth Syst.* **6,** 59–74. ISSN: 19422466. (2020) (Mar. 2014).

18. Hewitt, H., Fox-Kemper, B., Pearson, B., Roberts, M. & Klocke, D. The Small Scales of the Ocean May Hold the Key to Surprises. *Nat. Clim. Chang.* **12,** 496–499. ISSN: 1758-6798. (2023) (June 2022).

19. Girolamo, S. D., Schmid, P., Schulthess, T. & Hoefler, T. *SimFS: A Simulation Data Virtualizing File System Interface* in *Proceedings of the 33st IEEE International Parallel & Distributed Processing Symposium (IPDPS'19)* (IEEE, Rio de Janeiro, Brazil, 2019).

20. Becker, T. & Wing, A. A. Understanding the extreme spread in climate sensitivity within the radiative-convective equilibrium model intercomparison project. *Journal of Advances in Modeling Earth Systems* **12,** e2020MS002165 (2020).

21. Bony, S, Semie, A, *et al.* Observed modulation of the tropical radiation budget by deep convective organization and lower-tropospheric stability. *AGU advances* **1,** e2019AV000155 (2020).

22. Semie, A. G. & Bony, S. Relationship between precipitation extremes and convective organization inferred from satellite observations. *Geophysical Research Letters* **47,** e2019GL086927 (2020).

23. Roca, R. & Fiolleau, T. Extreme precipitation in the tropics is closely associated with long-lived convective systems. *Communications Earth & Environment* **1,** 18 (2020).

24. Hohenegger, C., Brockhaus, P., Bretherton, C. S. & Schär, C. The Soil Moisture–Precipitation Feedback in Simulations with Explicit and Parameterized Convection. *Journal of Climate* **22,** 5003–5020. ISSN: 1520-0442, 0894-8755. (2020) (Oct. 2009).

25. Barlage, M., Chen, F., Rasmussen, R., Zhang, Z. & Miguez-Macho, G. The importance of scale-dependent groundwater processes in land-atmosphere interactions over the central United States. *Geophysical Research Letters* **48,** e2020GL092171 (2021).

26. Yeager, S., Danabasoglu, G. & Chang, P. Exceptional multi-year prediction skill of the Kuroshio Extension in a high-resolution decadal prediction system (2023).

27. Reynolds, C., Williams, K. D. & Zadra, A. *WGNE systematic error survey results summary* in *100th American Meteorological Society Annual Meeting* (2020).

28. Bock, L., Lauer, A., *et al.* Quantifying Progress Across Different CMIP Phases With the ESMValTool. *Journal of Geophysical Research: Atmospheres* **125,** e2019JD032321. ISSN: 2169-8996. (2023) (2020).

29. Schiemann, R., Athanasiadis, P., *et al.* Northern Hemisphere Blocking Simulation in Current Climate Models: Evaluating Progress from the Climate Model Intercomparison Project Phase 5 to 6 and Sensitivity to Resolution. *Weather and Climate Dynamics* **1,** 277–292. (2023) (June 2020).

30. Moreno-Chamarro, E., Caron, L.-P., Ortega, P., Tomas, S. L. & Roberts, M. J. Can We Trust CMIP5/6 Future Projections of European Winter Precipitation? *Environ. Res. Lett.* **16,** 054063. ISSN: 1748-9326. (2023) (May 2021).

31. Moreno-Chamarro, E., Caron, L.-P., *et al.* Impact of Increased Resolution on Long-Standing Biases in HighResMIP-PRIMAVERA Climate Models. *Geoscientific Model Development* **15,** 269–289. ISSN: 1991-959X. (2023) (Jan. 2022).

32. Murakami, H., Vecchi, G. A., *et al.* Seasonal Forecasts of Major Hurricanes and Landfalling Tropical Cyclones Using a High-Resolution GFDL Coupled Climate Model. *Journal of Climate* **29,** 7977–7989. ISSN: 0894-8755, 1520-0442. (2023) (Nov. 2016).

33. Takasuka, D., Kodama, C., *et al. How Can We Improve the Seamless Representation of Climatological Statistics and Weather Toward Reliable Global K-scale Climate Simulations?* Preprint (Preprints, Mar. 2023). (2023).

34. Klocke, D., Brueck, M., Hohenegger, C. & Stevens, B. Rediscovery of the doldrums in storm-resolving simulations over the tropical Atlantic. *Nature Geoscience* **10,** 891–896 (2017).

35. Judt, F. & Rios-Berrios, R. Resolved convection improves the representation of equatorial waves and tropical rainfall variability in a global nonhydrostatic model. *Geophysical Research Letters* **48,** e2021GL093265 (2021).

36. Schemm, S. Toward Eliminating the Decades-Old "Too Zonal and Too Equatorward" Storm-Track Bias in Climate Models. *Journal of Advances in Modeling Earth Systems* **15,** e2022MS003482. ISSN: 1942-2466. (2023) (2023).

37. Sandu, I., van Niekerk, A., *et al.* Impacts of orography on large-scale atmospheric circulation. *npj Climate and Atmospheric Science* **2,** 10 (2019).

38. Hart, N. C., Washington, R. & Stratton, R. A. Stronger local overturning in convective-permitting regional climate model improves simulation of the subtropical annual cycle. *Geophysical Research Letters* **45,** 11–334 (2018).

39. Prein, A. F., Langhans, W., *et al.* A review on regional convection-permitting climate modeling: Demonstrations, prospects, and challenges. *Reviews of geophysics* **53,** 323–361 (2015).

40. Patricola, C. M. & Wehner, M. F. Anthropogenic influences on major tropical cyclone events. *Nature* **563,** 339–346 (2018).

41. Fowler, H. J., Lenderink, G., *et al.* Anthropogenic intensification of short-duration rainfall extremes. *Nature Reviews Earth & Environment* **2,** 107–122 (2021).

42. Cheng, K.-Y., Harris, L., *et al.* Impact of Warmer Sea Surface Temperature on the Global Pattern of Intense Convection: Insights From a Global Storm Resolving Model. *Geophysical Research Letters* **49,** e2022GL099796. ISSN: 1944-8007. (2023) (2022).

43. Rasmussen, R., Liu, C., *et al.* High-resolution coupled climate runoff simulations of seasonal snowfall over Colorado: a process study of current and warmer climate. *Journal of Climate* **24,** 3015–3048 (2011).

44. Kawase, H., Yamazaki, T., *et al.* Changes in extremely heavy and light snow-cover winters due to global warming over high mountainous areas in central Japan. *Progress in Earth and Planetary Science* **7,** 1–17 (2020).

45. Eidhammer, T., Booth, A., *et al.* Mass balance and hydrological modeling of the Hardangerjøkulen ice cap in south-central Norway. *Hydrology and Earth System Sciences* **25,** 4275–4297 (2021).

46. Hartick, C., Furusho-Percot, C., Clark, M. P. & Kollet, S. An Interannual Drought Feedback Loop Affects the Surface Energy Balance and Cloud Properties. *Geophysical Research Letters* **49,** e2022GL100924. ISSN: 1944-8007. (2023) (2022).

47. Doddridge, E. W., Marshall, J., *et al.* Eddy Compensation Dampens Southern Ocean Sea Surface Temperature Response to Westerly Wind Trends. *Geophys. Res. Lett.* **46,** 4365–4377. ISSN: 0094-8276, 1944-8007. (2020) (Apr. 2019).

48. Rackow, T., Danilov, S., *et al.* Delayed Antarctic Sea-Ice Decline in High-Resolution Climate Change Simulations. *Nat Commun* **13,** 637. ISSN: 2041-1723. (2023) (Feb. 2022).

49. Bauer, J. E., Cai, W.-J., *et al.* The Changing Carbon Cycle of the Coastal Ocean. *Nature* **504,** 61–70. ISSN: 1476-4687. (2023) (Dec. 2013).

50. Castleman, B. A., Schlegel, N.-J., Caron, L., Larour, E. & Khazendar, A. *Derivation of Bedrock Topography Measurement Requirements for the Reduction of Uncertainty in Ice Sheet Model Projections of Thwaites Glacier* Preprint (Ice sheets/Numerical Modelling, Oct. 2021). (2023).

51. Clerc, F., Minchew, B. M. & Behn, M. D. Marine Ice Cliff Instability Mitigated by Slow Removal of Ice Shelves. *Geophysical Research Letters* **46,** 12108–12116. ISSN: 1944-8007. (2023) (2019).

52. Jacob, C., Gettelman, A., Hohenegger, C., Senior, C. & Prein, A. *Report of the WCRP km-scale modeling workshop* tech. rep. (WCRP, 2022).

53. Lopez-Gomez, I., Cohen, Y., He, J., Jaruga, A. & Schneider, T. A Generalized Mixing Length Closure for Eddy-Diffusivity Mass-Flux Schemes of Turbulence and Convection. *J. Adv. Model. Earth Syst.* **12.** ISSN: 1942-2466, 1942-2466. (2023) (Nov. 2020).

54. Segura, H., Hohenegger, C., Wengel, C. & Stevens, B. Learning by Doing: Seasonal and Diurnal Features of Tropical Precipitation in a Global-Coupled Storm-Resolving Model. *Geophysical Research Letters* **49,** e2022GL101796. ISSN: 1944-8007. (2022) (2022).

55. Lopez-Gomez, I., Cohen, Y., He, J., Jaruga, A. & Schneider, T. A generalized mixing length closure for eddy-diffusivity mass-flux schemes of turbulence and convection. *Journal of Advances in Modeling Earth Systems* **12,** e2020MS002161 (2020).

56. Cohen, Y., Lopez-Gomez, I., *et al.* Unified entrainment and detrainment closures for extended eddy-diffusivity mass-flux schemes. *Journal of Advances in Modeling Earth Systems* **12,** e2020MS002162 (2020).

57. Grönquist, P., Yao, C., *et al.* Deep Learning for Post-Processing Ensemble Weather Forecasts. *Philosophical Transactions of the Royal Society A* **379** (2021).

58. Spiga, A., Hinson, D. P., *et al.* Snow precipitation on Mars driven by cloud-induced nighttime convection. *Nature Geoscience* **10,** 652–657 (2017).

59. Lefèvre, M., Spiga, A. & Lebonnois, S. Mesoscale modeling of Venus' bow-shape waves. *Icarus* **335,** 113376 (2020).

60. Bardet, D., Spiga, A. & Guerlet, S. Joint evolution of equatorial oscillation and interhemispheric circulation in Saturn's stratosphere. *Nature Astronomy* **6,** 804–811 (2022).

61. Wilkinson, M. D., Dumontier, M., *et al.* The FAIR Guiding Principles for Scientific Data Management and Stewardship. *Sci Data* **3,** 160018. ISSN: 2052-4463. (2023) (Mar. 2016).

62. Tuia, D., Schindler, K., *et al.* *Artificial Intelligence to Advance Earth Observation: A Perspective* May 2023. arXiv: `2305.08413 [cs, eess, stat]`. (2023).

63. Loescher, H. W., Vargas, R., *et al.* Building a Global Ecosystem Research Infrastructure to Address Global Grand Challenges for Macrosystem Ecology. *Earth's Future* **10,** e2020EF001696. ISSN: 2328-4277. (2023) (2022).

64. Kulmala, M. Build a Global Earth Observatory. *Nature* **553,** 21–23. (2023) (Jan. 2018).