

Advanced Computing for Climate Modelling: Integrating AI, High-Performance Computing and Earth-System Science

Miska Luoto

Professor, University of Helsinki, Finland

Abstract

Climate modelling requires the numerical representation of interacting atmospheric, oceanic, terrestrial, cryospheric, and biogeochemical processes across spatial and temporal scales. Increasing model resolution, process complexity, ensemble size, and data volume has made climate science one of the most computationally demanding areas of research. High-performance computing enables large coupled simulations, but higher computational capacity alone cannot resolve persistent problems associated with subgrid processes, structural uncertainty, parameter calibration, data assimilation, and computational energy consumption. Artificial intelligence offers complementary capabilities for learning unresolved processes, constructing emulators, improving parameter estimation, detecting model biases, downscaling coarse outputs, and accelerating data analysis.

This paper examines an integrated AI–high-performance-computing framework for advanced climate modelling. A structured integrative review is combined with a transparent conceptual simulation comparing conventional physics-based and AI–HPC hybrid workflows across six relative computing budgets. The modeled multi-variable error values are illustrative and do not represent observations, operational forecasts, or real climate projections. The simulation suggests that both workflows improve with additional computing resources, but the assumed hybrid workflow achieves a larger reduction in normalized error because it combines physical simulation with data-driven parameterization and computational acceleration.

The analysis emphasizes that machine learning should complement rather than replace physical understanding. Models trained on historical or present-day conditions may fail under unfamiliar climate regimes, while apparently accurate offline parameterizations can produce unstable coupled simulations. Reliable integration therefore requires conservation constraints, uncertainty quantification, independent climate-regime testing, reproducible software, transparent data provenance, and continuous comparison with observations. Other challenges include exascale software adaptation, data-transfer bottlenecks, numerical precision, unequal access to computing infrastructure, and the environmental footprint of large-scale computation. The paper concludes that the most credible pathway is a hybrid climate-modelling architecture in which physical laws provide structure, AI supports selected computational tasks, high-performance computing enables scale, and Earth-system science governs interpretation.

Keywords: artificial intelligence; climate simulation; Earth-system modelling; high-performance computing; machine learning; numerical modelling; scientific computing; uncertainty quantification.

1. Introduction

Climate is produced by interactions among the atmosphere, oceans, land, cryosphere, biosphere, and human activities. These components exchange energy, water, momentum, carbon, nutrients, aerosols, and other chemical constituents. Their interactions operate across scales ranging from microscopic cloud processes to global ocean circulation and from seconds to centuries. Climate modelling therefore requires both scientific representations of Earth-system processes and computational methods capable of solving large systems of equations.

General circulation models numerically represent atmospheric and oceanic motion using physical laws governing fluid dynamics, thermodynamics, radiation, and mass conservation. Earth-system models extend these foundations by incorporating processes such as vegetation dynamics, terrestrial carbon cycling, atmospheric chemistry, marine ecosystems, sea ice, ice sheets, and interactions between climate and the carbon cycle. Model components communicate through coupling systems that exchange variables while maintaining numerical and physical consistency.

Spatial resolution strongly affects what a model can represent. A global model divides Earth into three-dimensional grid cells. Processes larger than the cells may be resolved directly, while smaller processes must be approximated through parameterizations. Clouds, convection, turbulence, land-surface heterogeneity, and small ocean eddies are important examples. Increasing resolution can improve the explicit representation of some processes, but computational cost rises rapidly because the number of grid cells increases and shorter numerical time steps may become necessary.

Climate modelling also depends on ensembles. Because the climate system is nonlinear and affected by uncertain initial conditions, forcing pathways, parameters, and model structures, a single simulation cannot adequately represent the range of plausible outcomes. Ensemble experiments vary these conditions and evaluate the resulting distributions. Increasing ensemble size improves uncertainty analysis but substantially increases computational demand.

High-performance computing has enabled progressively more complex and higher-resolution simulations. Modern systems combine large numbers of central processing units, graphics processing units, accelerators, high-speed interconnects, parallel storage, and specialized software. Climate codes must divide calculations across many computing elements while minimizing communication delays and load imbalance.

The growth of computational capacity has not eliminated modelling uncertainty. Some processes remain too small or complex to resolve, while increasing resolution may reveal new interactions that require additional scientific treatment. Model development is therefore not a simple progression in which greater computing power automatically produces correct predictions.

Artificial intelligence introduces complementary methods. Neural networks, random forests, Gaussian processes, and other machine-learning approaches can learn relationships from observations or high-resolution simulations. Potential applications include parameterizing unresolved processes, emulating expensive model components, calibrating parameters, detecting systematic biases, assimilating observations, statistically downscaling coarse outputs, and analyzing extensive simulation archives.

AI also introduces new risks. A model can fit historical data without learning physically valid relationships. It may violate conservation laws, perform poorly outside its training distribution, or destabilize a coupled climate simulation. A fast surrogate model can reproduce mean behavior while failing to represent extremes or feedbacks. Consequently, predictive accuracy on a held-out dataset is not sufficient evidence of climate-model reliability.

The central research problem is how AI, high-performance computing, and Earth-system science can be integrated without weakening physical consistency, uncertainty transparency, or scientific interpretability. This paper examines that problem through an integrative review and a conceptual simulation. The simulated results illustrate computational relationships and are not empirical climate findings.

2. Review of Literature

2.1 Evolution of Earth-System Modelling and Computational Scale

Climate models have developed from simplified representations of atmospheric energy balance into coupled systems containing atmosphere, ocean, land, sea ice, chemistry, and biogeochemical processes. Edwards described climate models as scientific and computational constructions shaped by observations, theory, computing capacity, and institutional coordination. Their credibility emerges from physical formulation, evaluation, intercomparison, and the ability to explain observed patterns rather than from computational size alone.

Coupled Model Intercomparison Projects established common experimental protocols for comparing climate models. Taylor et al. described the experimental design of a major intercomparison phase that supported systematic analysis of historical simulations, future scenarios, and model differences. Eyring et al. later presented a subsequent intercomparison framework incorporating expanded scientific questions and more coordinated experiments.

Model intercomparison reveals both robustness and persistent disagreement. Models may reproduce major features of the global climate while differing in regional precipitation, cloud feedback, ocean circulation, land-atmosphere interactions, or extremes. Ensemble analysis should therefore distinguish between uncertainty caused by initial conditions, model structure, parameters, and external forcing assumptions.

High-performance computing supports finer grids, larger ensembles, more complete process coupling, and longer simulations. However, parallel performance becomes increasingly difficult as models scale across large machines. Calculations involving local grid cells may parallelize effectively, while coupling and communication across components can create bottlenecks. Input/output operations also become expensive when simulations produce petabyte-scale archives.

Computational design is becoming heterogeneous. Climate models originally developed for conventional processors must increasingly operate across graphics processors and specialized accelerators. Adapting legacy scientific codes requires changes to algorithms, data structures, memory access, numerical precision, and software engineering practices. Scientific reproducibility must be maintained during these transitions.

2.2 Artificial Intelligence for Climate Processes and Model Acceleration

Machine learning has become increasingly relevant to Earth-system science because the field contains large observational archives, complex nonlinear relationships, and expensive simulations. Reichstein et al. argued that deep learning could contribute to Earth-system understanding when combined with process knowledge rather than treated as a purely predictive instrument. This combination is particularly important when observations are sparse, causal relationships are uncertain, or future conditions differ from training data.

One major application is parameterization. Rasp et al. trained a deep neural network to represent subgrid processes using high-resolution simulation information. Their work demonstrated that a learned parameterization could be coupled with a simplified atmospheric model and reproduce several statistical characteristics of the reference simulation. Gentine et al. similarly examined deep learning for representing cloud-related subgrid processes.

Brenowitz and Bretherton showed that data-driven parameterizations can encounter stability problems when integrated into prognostic models. A machine-learning model may perform well when tested offline against stored data yet produce unrealistic feedbacks during long coupled simulation. This distinction between offline accuracy and online stability is central to responsible climate-AI research. Yuval and O’Gorman developed a machine-learning parameterization designed to maintain stability across different model resolutions. Their work illustrates the importance of scale awareness and physically meaningful input–output relationships. O’Gorman and Dwyer investigated random forests for convection parameterization, providing an interpretable alternative to some neural-network approaches.

AI can also support emulation. An emulator approximates the input–output behavior of an expensive model component or complete simulation. Emulators may accelerate uncertainty analysis, scenario exploration, parameter calibration, and sensitivity testing. Their reliability depends on the coverage of training experiments. Extrapolation beyond the simulated parameter or climate range remains risky.

2.3 Data Assimilation, Downscaling and Model Evaluation

Data assimilation combines observations with model states to estimate the evolving condition of the Earth system. Numerical weather prediction has a long history of variational and ensemble-based data assimilation, while climate applications use related methods for initialization, reanalysis, and state estimation. Machine learning may help approximate observation operators, estimate model error, identify nonlinear relationships, and accelerate computationally expensive assimilation steps.

Downscaling translates coarse global information into finer regional detail. Dynamical downscaling uses regional climate models and preserves a strong physical foundation, but it is computationally demanding. Statistical downscaling learns relationships between large-scale climate variables and local outcomes. Machine learning expands this capacity through nonlinear models, spatial networks, and generative approaches. Downscaled products must preserve physical consistency and should not present artificial fine-scale detail as independently simulated information.

Model evaluation compares simulations with observations and reanalyses. Eyring et al. presented community-oriented evaluation tools capable of calculating standardized diagnostics across multiple models. Machine learning can detect complex bias patterns, cluster model behavior, and identify

relationships among variables. It should supplement rather than replace established process-based evaluation.

Uncertainty quantification remains essential. Climate models do not produce exact descriptions of future conditions. They generate conditional projections based on physical models, initial states, parameters, and external forcing scenarios. AI components introduce additional uncertainty related to training data, architecture, optimization, and distribution shift.

Table 1: Core functions within an integrated climate-computing architecture

System component	Principal role	Potential AI contribution	Major scientific risk
Dynamical core	Solves atmospheric or oceanic motion	Solver acceleration and learned numerical correction	Loss of stability or conservation
Physical parameterizations	Represent unresolved clouds, convection, turbulence, and radiation	Data-driven or hybrid parameterization	Poor performance outside training regimes
Land and biogeochemical models	Simulate water, vegetation, carbon, and nutrient processes	Surrogate models and parameter estimation	Oversimplification of ecological feedbacks
Data assimilation	Combines observations with model states	Learned operators and error estimation	Biased observations may influence state estimates
Downscaling	Produces regional or local information	Nonlinear statistical refinement	Artificial detail and unreliable extremes
Model evaluation	Diagnoses biases and process relationships	Pattern recognition and automated diagnostics	High predictive scores may conceal physical errors

The literature suggests that AI is most valuable when applied selectively to identifiable computational or scientific bottlenecks. Fully replacing process-based climate models with unconstrained statistical systems would weaken the ability to reason about physical mechanisms and unprecedented conditions.

3. Objectives

The principal objective of this study is to evaluate how artificial intelligence and high-performance computing can be integrated with Earth-system science to improve the efficiency, resolution, and

analytical capability of climate modelling. The paper focuses on hybrid systems that preserve physical laws while applying machine learning to selected processes.

A second objective is to develop evaluation principles that prevent computational acceleration from being mistaken for scientific improvement. A faster model is useful only when its accuracy, stability, uncertainty, and physical consistency remain adequate for the intended research question.

The specific objectives are:

- To explain the computational structure and scientific role of coupled Earth-system models.
- To examine how high-performance computing supports spatial resolution, ensemble simulation, process coupling, and climate-data analysis.
- To evaluate machine-learning applications in parameterization, emulation, downscaling, data assimilation, and model diagnosis.
- To distinguish offline machine-learning accuracy from stable performance inside a coupled climate model.
- To demonstrate an illustrative comparison between conventional and AI-HPC hybrid workflows across increasing computing budgets.
- To identify scientific, computational, environmental, institutional, and ethical barriers to advanced climate modelling.
- To propose validation requirements for physics-informed and data-driven climate components.
- To clarify the continuing importance of human expertise, observations, and process understanding.

The paper addresses five research questions:

1. How can AI improve the computational efficiency and representation of unresolved processes in climate models?
2. Which high-performance-computing capabilities are most important for next-generation Earth-system simulation?
3. How should hybrid climate models be evaluated for accuracy, stability, conservation, and transferability?
4. What scientific and institutional risks arise from increasing dependence on AI and large-scale computing?
5. How can advanced computing support trustworthy and socially useful climate information?

4. Research Methodology

4.1 Evidence Identification and Eligibility

A structured integrative-review approach was employed. Scholarly evidence was identified from multidisciplinary sources covering climate dynamics, Earth-system science, atmospheric modelling, ocean modelling, machine learning, scientific computing, data assimilation, and uncertainty quantification. Information sources included Web of Science-oriented publications, Scopus-oriented academic records, Google Scholar, publisher databases, the American Geophysical Union, Nature Portfolio, the National Academy of Sciences, and Geoscientific Model Development.

The final literature search was conducted on August 25, 2026, while references included in the manuscript were restricted to the stated bibliographic boundary. Search combinations included machine learning climate model parameterization, high-performance computing Earth-system model, deep learning subgrid climate processes, climate model emulation, AI climate downscaling, data assimilation machine learning, climate uncertainty quantification, and physics-informed climate modelling.

Studies were eligible when they addressed Earth-system simulation, climate-model development, computational scaling, data-driven parameterization, model emulation, downscaling, data assimilation, or climate-model evaluation. Priority was given to peer-reviewed studies with identifiable methods and direct relevance to climate modelling.

Evidence focused only on short-term weather forecasting was included selectively when it addressed methods transferable to climate modelling. Promotional computing claims, undocumented benchmarks, and unsupported performance assertions were excluded from the analytical core.

4.2 Analytical Framework and Evaluation Criteria

The evidence was coded across six dimensions: physical process represented, computational role, machine-learning method, training-data source, validation strategy, and reported limitation. A second coding layer examined reproducibility, uncertainty reporting, conservation constraints, computational cost, and testing under unfamiliar climate conditions.

A normalized multi-variable error index can be represented conceptually as:

$$E = v \cdot 1 \sum n w v (\sigma v RMSE v) + \lambda C + \gamma S$$

where:

- E represents the integrated model-error index;
- RMSEv is the root-mean-square error for climate variable v;
- σv is a normalization scale;
- wv is the importance assigned to variable v;
- C is a penalty for conservation violations;
- S is a penalty for numerical instability; and
- λ and γ are penalty weights.

This formulation is conceptual. A real climate evaluation would require region-specific, variable-specific, seasonal, process-oriented, and extreme-event diagnostics rather than a single aggregated score.

4.3 Conceptual Simulation Design

A deterministic illustrative simulation was constructed to compare conventional physics-based and AI-HPC hybrid climate-modelling workflows. Six relative compute budgets were used: 1, 2, 4, 8, 16, and 32 times an arbitrary baseline.

The conventional workflow assumes that additional resources are directed toward resolution, ensemble size, and process complexity. The hybrid workflow assumes that the same computational growth is combined with learned parameterizations, surrogate components, workflow optimization, and accelerated analysis. The values were chosen for conceptual illustration and were not derived from an operational climate model.

Table 2: Supporting data for the simulated computing comparison

Relative compute budget	Conventional error index	AI-HPC hybrid error index	Hybrid reduction relative to conventional
1×	3.20	3.05	4.7%
2×	2.92	2.62	10.3%
4×	2.66	2.21	16.9%
8×	2.47	1.86	24.7%
16×	2.33	1.58	32.2%
32×	2.24	1.39	37.9%

The error index is dimensionless, and lower values indicate better modeled performance under the assumed criteria. No inferential statistics were applied because the numbers are simulation inputs rather than sampled observations.

5. Analysis

Climate modelling is constrained by a three-way interaction among physical complexity, spatial resolution, and ensemble size. Increasing one dimension may limit the others under a fixed computing budget. A very high-resolution simulation may permit only a small ensemble, making uncertainty characterization difficult. A large ensemble at coarse resolution may quantify variability but miss important regional processes.

Advanced computing should therefore be allocated according to the scientific question. Global feedback analysis may require long coupled simulations and diverse forcing scenarios. Regional extreme-event research may place greater value on spatial resolution. Carbon-cycle research may prioritize process coupling and biogeochemical complexity.

High-performance computing enables domain decomposition, in which the model grid is divided across many processing units. Effective scaling requires balanced workloads and efficient communication. Atmosphere, ocean, land, and sea-ice components may have different computational characteristics, complicating resource allocation.

Graphics processors offer high throughput for suitable numerical operations, particularly dense linear algebra and machine-learning training. Climate codes developed over decades may not map efficiently onto these architectures. Portability frameworks and redesigned algorithms can reduce dependence on a single machine, but they require sustained software-engineering investment.

Memory and data movement are increasingly important constraints. Moving information between processors or writing large output archives may consume more time and energy than arithmetic operations. In-situ analysis can reduce this burden by processing data while the simulation is running and storing only selected diagnostics.


AI can provide value through learned parameterization. A high-resolution simulation can generate training data for a neural network representing unresolved cloud or convective behavior in a coarser model. The learned component may be much faster than the original high-resolution calculation.

The approach is scientifically credible only if the learned component remains stable after coupling. Offline validation tests predictions against stored reference values. Online validation evaluates the component during a long simulation in which its outputs influence future model states. Small systematic errors may accumulate or trigger unrealistic feedbacks.

Physical constraints can be introduced through architecture design, loss functions, output corrections, or hybrid formulations. Conservation of mass, energy, water, and relevant chemical constituents should be preserved to an appropriate numerical tolerance. Non-negativity constraints may be needed for quantities such as humidity or tracer concentrations.

Emulators provide another form of acceleration. A trained surrogate can approximate an expensive model's response across a defined parameter space. Researchers can use it for sensitivity analysis, calibration, and uncertainty propagation. The emulator should not be applied beyond the range of training experiments without explicit extrapolation analysis.

Graph Title: Simulated Climate-Model Error Across Relative Compute Budgets
Simulated Climate-Model Error Across Relative Compute Budgets



Both simulated workflows show decreasing error as computing resources increase. The conventional workflow declines from 3.20 to 2.24, representing a 30% reduction. The AI-HPC hybrid workflow declines from 3.05 to 1.39, representing an approximate 54.4% reduction.

6. Challenges

Climate models are approximations of a complex system. Even high-resolution simulations cannot represent every relevant process explicitly. Structural uncertainty remains when processes are simplified, omitted, or represented differently among models.

Machine-learning systems depend on training data. Observational datasets contain gaps, instrument changes, spatial inequalities, and measurement uncertainty. Simulation-generated training data inherit the assumptions and biases of the reference model. An AI component trained on biased data may reproduce those biases efficiently.

Distribution shift is particularly important for climate research. The purpose of climate projection is to examine conditions that may differ from the historical record. A model trained on present-day relationships may encounter temperatures, moisture conditions, circulation regimes, or ecosystem responses outside its training distribution.

Conservation violations can create severe problems. Small imbalances in energy or water may accumulate during long simulations. Physics-aware learning and explicit correction can reduce this risk, but constraint enforcement may also introduce new numerical interactions requiring evaluation.

Hybrid instability is a major barrier. A learned parameterization may be accurate on independent samples while producing unstable coupled behavior. Online testing is therefore essential before replacing established model components.

Interpretability remains limited. A neural network may detect useful relationships without providing a process explanation. Feature-attribution methods can help, but they do not establish causality. Climate scientists should examine whether learned relationships remain physically plausible under changing conditions.

Model calibration may improve agreement with selected observations while degrading unmeasured processes or regions. Calibration objectives must balance multiple variables, seasons, locations, and process relationships. Excessive tuning to the historical climate may reduce future reliability.

High-performance-computing systems create software challenges. Codes must function across changing processor architectures, compilers, libraries, and storage systems. Scientific teams require long-term support for software testing, documentation, version control, and reproducible workflows.

Numerical precision creates additional tradeoffs. Reduced precision can accelerate machine-learning and selected simulation calculations, but rounding errors may accumulate in sensitive components. Mixed-precision strategies require component-specific validation.

Data storage is another constraint. Saving every variable at every model level and time step is impractical. Researchers must decide what to retain before all future analyses are known. In-situ diagnostics and standardized archives can help, but overly aggressive reduction may eliminate valuable information.

Reproducibility can be difficult on large machines. Hardware, parallel decomposition, compiler options, and nondeterministic accelerator operations may influence results. Exact bitwise reproduction may not always be feasible, but scientific conclusions should remain robust.

Unequal access to computing infrastructure may deepen disparities in climate research. Institutions in climate-vulnerable regions may lack access to supercomputers, extensive observations, storage, and technical personnel.

This can limit local control over modelling priorities and regional interpretation.

Open-source models and shared data can reduce barriers, but large datasets remain difficult to transfer and analyze. Capacity building should include software engineering, climate science, data management, and access to suitable computational resources.

Communication presents a final challenge. High-resolution visual outputs may appear precise even when underlying uncertainties remain large. Climate information should clearly distinguish resolution, accuracy, likelihood, scenario dependence, and model agreement.

7. Discussion

7.1 Hybrid Modelling as a Scientific Architecture

The most credible role for AI is within hybrid architectures where physical equations govern resolved dynamics and machine learning addresses selected computational bottlenecks. This preserves scientific structure while allowing data-driven improvement.

Hybridization should be selective. A mature physical parameterization should not be replaced merely because a neural network achieves lower offline error. Replacement is justified when the learned component provides demonstrable gains in coupled stability, process representation, computational efficiency, and uncertainty characterization.

The design should also preserve reversibility. Researchers should be able to compare simulations with and without the learned component, examine failure cases, and restore a validated conventional option. Modular implementation supports such comparisons.

7.2 From Computational Performance to Climate Credibility

Faster execution and lower benchmark error are necessary but insufficient indicators. Climate credibility depends on conservation, long-term stability, regional behavior, variability, extremes, feedbacks, and performance under unfamiliar conditions.

Evaluation should include process-oriented diagnostics. A model may reproduce precipitation averages for the wrong physical reason, creating compensating errors. Examining energy budgets, moisture transport, circulation patterns, cloud regimes, ocean heat uptake, and carbon feedbacks helps determine whether improved output reflects better process representation.

Uncertainty must remain visible. AI systems should not collapse ensembles into a single apparently precise outcome. Their role should include efficient generation, calibration, and interpretation of probabilistic information.

7.3 Infrastructure, Equity and Environmental Responsibility

Advanced climate modelling depends on substantial public and institutional infrastructure. Supercomputers, observation networks, data archives, software teams, and international model intercomparison efforts all contribute to model credibility.

The distribution of these resources affects whose questions are investigated. Regional researchers should participate in model design, evaluation, and interpretation rather than being treated only as recipients of global outputs. Shared computing services, regional training, open software, and accessible diagnostics can support greater participation.

Computational environmental impact should be considered alongside scientific benefit. Efficient algorithms, renewable-powered facilities, transparent energy reporting, and the reuse of simulation outputs can reduce avoidable burdens. The objective is not minimal computation but scientifically justified and responsibly managed computation.

8. Conclusion

Advanced climate modelling is entering a period in which physical simulation, artificial intelligence, and high-performance computing are increasingly interdependent. Earth-system models provide a physically structured account of interacting climate components, while HPC enables greater resolution, longer simulations, and larger ensembles. AI can accelerate selected calculations and help represent complex unresolved processes.

The conceptual simulation illustrates a possible advantage of hybrid AI–HPC workflows across increasing computing budgets. Under the stated assumptions, the hybrid error index falls from 3.05 to 1.39 and shows a larger improvement than the conventional workflow. These values are illustrative and cannot be interpreted as evidence about an operational model or future climate projection.

The scientific value of hybrid systems depends on more than predictive accuracy. Learned components must be evaluated for online stability, physical consistency, climate-regime transferability, uncertainty calibration, resolution dependence, and computational cost. Independent observations and process-based diagnostics remain essential.

The strongest future pathway is not a choice between physics and AI. It is a carefully governed integration in which physical knowledge establishes constraints, AI supports targeted approximation and analysis, HPC provides computational scale, and Earth-system scientists retain responsibility for validation and interpretation. This architecture can improve climate information while preserving the transparency and caution required for high-consequence environmental decisions.

References

1. Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N., & Prabhat. (2019). Deep learning and process understanding for data-driven Earth system science. *Nature*, 566, 195–204. <https://doi.org/10.1038/s41586-019-0912-1>
2. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521, 436–444. <https://doi.org/10.1038/nature14539>
3. Rolnick, D., Donti, P. L., Kaack, L. H., Kochanski, K., Lacoste, A., Sankaran, K., Ross, A. S., Milojevic-Dupont, N., Jaques, N., Waldman-Brown, A., Luccioni, A. S., Maharaj, T., Sherwin, E. D., Mukkavilli, S. K., Kording, K. P., Gomes, C. P., Ng, A. Y., Hassabis, D., Platt, J. C., ... Bengio, Y. (2019). Tackling climate change with machine learning. *arXiv*. <https://arxiv.org/abs/1906.05433>

4. Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. *Nature Reviews Physics*, 3, 422–440. <https://doi.org/10.1038/s42254-021-00314-5>
5. Rasp, S., Pritchard, M. S., & Gentine, P. (2018). Deep learning to represent subgrid processes in climate models. *Proceedings of the National Academy of Sciences*, 115(39), 9684–9689. <https://doi.org/10.1073/pnas.1810286115>
6. Brenowitz, N. D., & Bretherton, C. S. (2018). Prognostic validation of a neural network unified physics parameterization. *Geophysical Research Letters*, 45(12), 6289–6298. <https://doi.org/10.1029/2018GL078510>
7. Gentine, P., Pritchard, M., Rasp, S., Reinaudi, G., & Yacalis, G. (2018). Could machine learning break the convection parameterization deadlock? *Geophysical Research Letters*, 45(11), 5742–5751. <https://doi.org/10.1029/2018GL078202>
8. Dueben, P. D., & Bauer, P. (2018). Challenges and design choices for global weather and climate models based on machine learning. *Geoscientific Model Development*, 11(10), 3999–4009. <https://doi.org/10.5194/gmd-11-3999-2018>
9. Ham, Y.-G., Kim, J.-H., & Luo, J.-J. (2019). Deep learning for multi-year ENSO forecasts. *Nature*, 573, 568–572. <https://doi.org/10.1038/s41586-019-1559-7>
10. Weyn, J. A., Durran, D. R., & Caruana, R. (2020). Improving data-driven global weather prediction using deep convolutional neural networks on a cubed sphere. *Journal of Advances in Modeling Earth Systems*, 12(9), e2020MS002109. <https://doi.org/10.1029/2020MS002109>
11. Sønderby, C. K., Espeholt, L., Heek, J., Dehghani, M., Oliver, A., Salimans, T., Agrawal, S., Hickey, J., & Kalchbrenner, N. (2020). MetNet: A neural weather model for precipitation forecasting. *arXiv*. <https://arxiv.org/abs/2003.12140>
12. Ravuri, S., Lenc, K., Willson, M., Kangin, D., Lam, R., Mirowski, P., Fitzsimons, M., Athanassiadou, M., Bhatia, K., Engin, O., Ruede, M., et al. (2021). Skilful precipitation nowcasting using deep generative models of radar. *Nature*, 597, 672–677. <https://doi.org/10.1038/s41586-021-03854-z>
13. Pathak, J., Subramanian, S., Harrington, P., Raja, S., Chattopadhyay, A., Mardani, M., Kurth, T., Hall, D., Li, K., Azizzadenesheli, K., & others. (2022). FourCastNet: A global data-driven high-resolution weather model using adaptive Fourier neural operators. *arXiv*. <https://arxiv.org/abs/2202.11214>
14. Bi, K., Xie, L., Zhang, H., Chen, X., Gu, X., & Tian, Q. (2023). Accurate medium-range global weather forecasting with 3D neural networks. *Nature*, 619, 533–538. <https://doi.org/10.1038/s41586-023-06185-3>
15. Lam, R., Sanchez-Gonzalez, A., Willson, M., Wirsberger, P., Fortunato, M., Alet, F., Ravuri, S., Ewalds, T., Eaton-Rosen, Z., Hu, W., Merose, A., Hoyer, S., Holland, G., Vinyals, O., Stott, J., Pritzel, A., Mohamed, S., & Battaglia, P. W. (2023). Learning skillful medium-range global weather forecasting. *Science*, 382(6677), 1416–1421. <https://doi.org/10.1126/science.adi2336>
16. Eyring, V., Bony, S., Meehl, G. A., Senior, C. A., Stevens, B., Stouffer, R. J., & Taylor, K. E. (2016). Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization. *Geoscientific Model Development*, 9(5), 1937–1958. <https://doi.org/10.5194/gmd-9-1937-2016>



International Journal of Innovative Computer Science and IT Research

Volume 2, Issue 9 – September 2026 • E-ISSN: 3067-1108 P-ISSN: 3067-1094

17. Taylor, K. E., Stouffer, R. J., & Meehl, G. A. (2012). An overview of CMIP5 and the experiment design. *Bulletin of the American Meteorological Society*, 93(4), 485–498. <https://doi.org/10.1175/BAMS-D-11-00094.1>
18. Dunne, J. P., Zhang, Y., Stouffer, R. J., John, J. G., Held, I. M., Dunne, K. A., ... & others. (2020). The GFDL Earth System Model version 4.1 (GFDL-ESM4.1): Physical formulation and simulation characteristics. *Journal of Advances in Modeling Earth Systems*, 12(10), e2019MS002015. <https://doi.org/10.1029/2019MS002015>
19. Balaji, V. (2015). Computational challenges in climate modelling. *Journal of Physics: Conference Series*, 640, 012014. <https://doi.org/10.1088/1742-6596/640/1/012014>
20. Bauer, P., Thorpe, A., & Brunet, G. (2015). The quiet revolution of numerical weather prediction. *Nature*, 525, 47–55. <https://doi.org/10.1038/nature14956>