

1 Machine Learning for numerical weather and climate modelling: 2 a review

3 Catherine O. de Burgh-Day¹ & Tennessee Leeuwenburg¹

4 ¹The Bureau of Meteorology, 700 Collins St Docklands, Victoria, Australia

5 Correspondence to: Catherine O. de Burgh-Day (catherine.deburgh-day@bom.gov.au)

6 Abstract.

7 Machine learning (ML) is increasing in popularity in the field of weather and climate modelling. Applications range
8 from improved solvers and preconditioners, to parameterization scheme emulation and replacement, and more recently
9 even to full ML-based weather and climate prediction models. While ML has been used in this space for more than
10 25 years, it is only in the last 10 or so years that progress has accelerated to the point that ML applications are becoming
11 competitive with numerical knowledge-based alternatives. In this review, we provide a roughly chronological
12 summary of the application of ML to aspects of weather and climate modelling from early publications through to the
13 latest progress at the time of writing. We also provide an overview of key ML terms, methodologies, and ethical
14 considerations. Finally, we discuss some potentially beneficial future research directions. Our aim is to provide a
15 primer for researchers and model developers to rapidly familiarize and update themselves with the world of ML in the
16 context of weather and climate models.

Deleted: parametrisation

Deleted: concepts

Deleted: and

Deleted: , the selection of ML techniques

Deleted: , and

17 1. Introduction

18 Current state-of-the-art weather and climate models use numerical methods to solve equations representing the
19 dynamics of the atmosphere and ocean on meshed grids. The grid-scale effects of processes that are too small to be
20 resolved are either represented by parametrization schemes or are prescribed. These numerical weather and climate
21 forecasts are computationally costly and are not easy to implement on specialized compute resources such as GPUs
22 (although there are efforts underway to do so, for example in LFRic (Adams et al. 2019)). One of the main approaches
23 to improving forecast accuracy is to increase model resolution (reduced timestep between model increments and/or
24 decreased grid spacing), but due to the high computational cost of this approach, improvements in model skill are
25 hampered by the finite supercomputer capacity available. An additional pathway to improve skill is to improve the
26 understanding and representation of subgrid-scale processes, however this is again a potentially computationally costly
27 exercise.

Deleted: amenable

Deleted: transfer to

Deleted: sub grid-scale

28 In the remainder of this introduction, we overview the state of machine learning in weather and climate research
29 without always providing references; we instead provide relevant references in the detailed sections that follow.

Deleted:

30 Machine learning is an increasingly powerful and popular tool. It has proven to be computationally efficient, as well
31 as being an accurate way to model subgrid-scale processes. The term “Machine learning” (ML) was first coined by

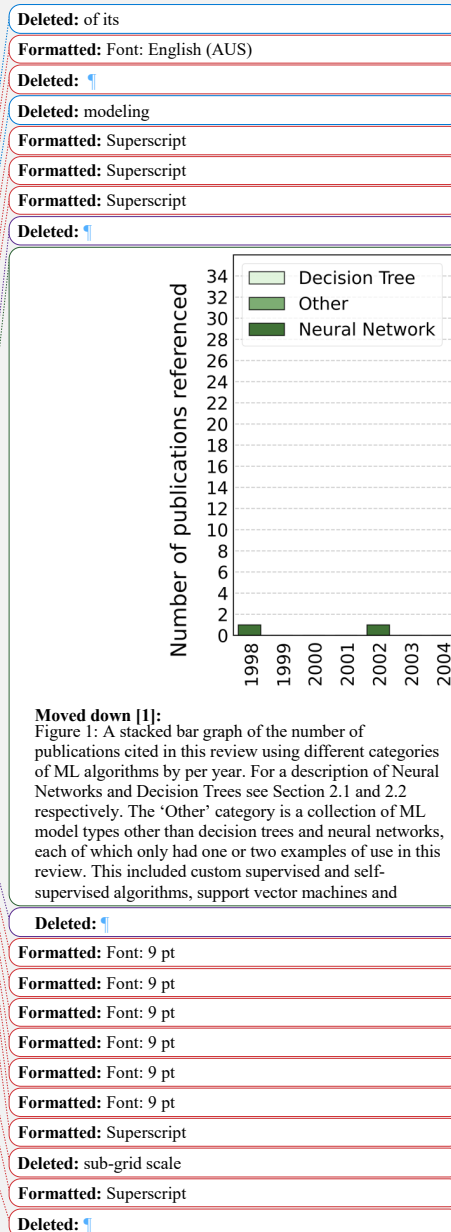
Deleted: sub-grid scale

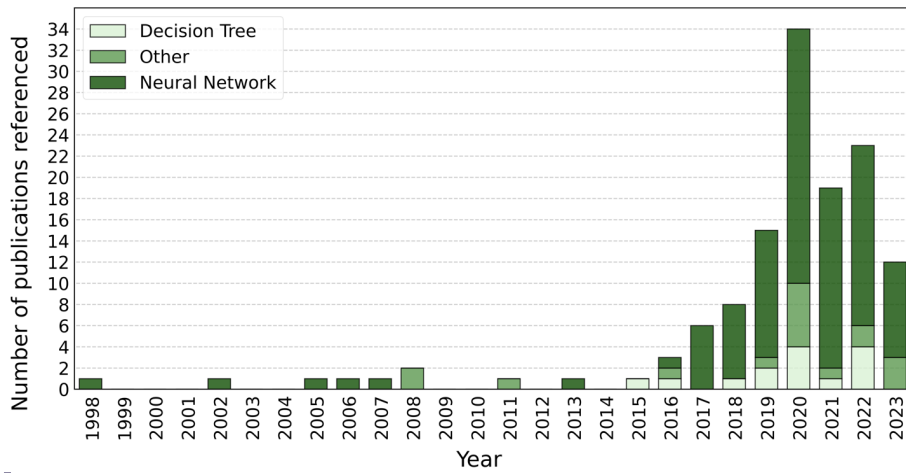
42 Arthur Samuel in 1952 to refer to a “field of study that gives computers the ability to learn without being explicitly
 43 programmed”¹. Learning by example is the defining characteristic of ML.
 44 The growing potential for ML in weather and climate modelling is being increasingly recognized by meteorological
 45 agencies and researchers around the world. The former is evidenced by the development of strategies and frameworks
 46 to better support the development of ML research, such as the Data Science Framework recently published by the Met
 47 Office in the UK². The latter is made clear by the explosion in publications from academia, government agencies and
 48 private industry in this space, as demonstrated by the rest of this review. **Figure 1 shows the number of publications**
 49 **cited in this review using different categories of ML algorithms by year, and clearly illustrates the increase in the**
 50 **uptake of ML methods by the research community.**
 51 **This is not necessarily an unbiased sample of the use of different architectures in the literature, since the selection of**
 52 **papers cited in this review focuses on telling the story of the growth of the use of ML in weather and climate modelling**
 53 **over time, rather than being a comprehensive list of all uses of ML in the literature.**
 54 There are established techniques and aspects of the weather and climate **modelling** lifecycle that would already be
 55 considered ML by many. For example, linear regression³, principal component analysis, correlations, and the
 56 calculation of teleconnections can all be considered types of ML. Data Assimilation techniques could also be
 57 considered a form of ML. There are, however, other classes of ML (e.g. Neural Networks[†], Decision Trees[†], etc.)
 58 which are much less widely used within the weather and climate modelling space and have great potential to be of
 59 benefit. There is growing interest in, and increasingly effective application of, these ML techniques to take the place
 60 of more traditional approaches to modelling. The potential for ML in weather and climate modelling extends all the
 61 way from replacement of individual sub-components of the model (to improve accuracy and reduce computational
 62 cost) to full replacement of the entire numerical model.
 63 **While ML models are typically computationally costly during training, they can provide very fast predictions at**
 64 **inference[†] time, especially on GPU hardware. They often also avoid the need to have full understanding of the**
 65 **processes being represented and can learn and infer complex relationships without any need for them to be explicitly**
 66 **encoded. These properties make ML an attractive alternative to traditional parametrization, numerical solver, and**
 67 **modelling methods.**
 68 Neural Networks (NNs, explained further in Section 2.1) in particular are an increasingly favored alternative approach
 69 for representing **sub-grid-scale** processes or replacing numerical models entirely. They consist of several
 70 interconnected layers of nonlinear nodes[†], with the number of intermediate layers depending on the complexity of the
 71 system being represented. These nodes allow for the encoding of an arbitrary number of interrelationships between
 72 arbitrary parameters to represent the system, removing the need to explicitly encode these interrelationships into a
 73 parameterization or numerical model.

¹ <http://infolab.stanford.edu/pub/voy/museum/samuel.html>, accessed 7th February 2023

² <https://www.metoffice.gov.uk/research/foundation/informatics-lab/met-office-data-science-framework>, accessed 7 February 2023

³ Henceforth, the first occurrence of each term described in the glossary is marked with the symbol "†"





Moved (insertion) [1]
 Deleted: <object>
 Formatted: Indent: Left: 0 cm

112
 113 Figure 1: A stacked bar graph of the number of publications cited in this review using different categories of ML algorithms by per
 114 year. For a description of Neural Networks and Decision Trees see Section 2.1 and 2.2 respectively. The ‘Other’ category is a
 115 collection of ML model types other than decision trees and neural networks, each of which only had one or two examples of use in
 116 this review. This included custom supervised and self-supervised algorithms, support vector machines and relevance vector
 117 machine models, regression models, unsupervised learning models, reservoir computing models and non-NN gaussian models.
 118 This figure includes all references from this review except for: seminal ML papers that are on new ML methods (e.g., foundational
 119 ML papers from outside the domain of weather and climate modelling), review papers, any paper cited that concerns a topic which
 120 is out of scope (e.g., nowcasting), and any other paper which does not present a new method directly applicable to weather and
 121 climate modelling. The full table of citations is provided in the appendix.
 122
 123

124 One challenge that must be overcome before there will be more widespread acceptance of ML as an alternative to
 125 traditional modelling methods is that ML is seen as lacking interpretability. Most ML models do not explicitly
 126 represent the physical processes they are simulating, although physics constrained ML is a new and growing field
 127 which goes some way to addressing this (see Section 6). Furthermore, the techniques available to gain insight into the
 128 relative importance and predictive mechanism of each predictor (i.e. the model outputs) are limited. In contrast,
 129 traditional models are usually driven by some understanding and/or representation of the physical mechanisms and
 130 processes which are occurring. This makes it possible to more easily gain insight into what physical drivers could
 131 explain a given output. The “black box” nature of many current ML approaches to parametrization makes them an
 132 unpopular choice for many researchers, (and can be off-putting for decision makers), since, for example, explaining
 133 what went wrong in a model after a bad forecast can be more challenging if there are processes in the model which
 134 are not, and cannot, be understood through the lens of physics. However, increasing attention is being paid to the
 135 interpretability of ML models (e.g., McGovern et al., 2019; Toms et al., 2020; Samek et al., 2021), and there are

Deleted: ,
 Deleted: ,

139 existing methods to provide greater insight into the way physical information is propagated through them (e.g.,
140 attention maps, which identify the regions in spatial input data that have the greatest impact on the output field, and
141 ablation studies, which involve comparing reduced data sources and/or models to the original models that have full
142 access to available data, to gain insight into the models).

143 As with their traditional counterparts, ML-based parametrizations and emulators are typically initially developed in
144 single-column models, aquaplanet configurations, or otherwise simplified models. There are many examples of ML-
145 based schemes which have been shown to perform well against benchmark alternatives in this setting, only to fail to
146 do so in a realistic model setting. A common theme is that these ML schemes rapidly excite instabilities in the model
147 as errors in the ML parametrization push key parameters outside of the domain of the training data as the overall
148 model is integrated forward in time, leading to rapidly escalating errors and to the model ‘blowing up’. Similarly,
149 many ML-based full model replacements perform well for short lead times, only to exhibit model drift and a rapid
150 loss of skill for longer lead times due to rapidly growing errors and the model drifting outside its training envelope.

151 In recent years, however, progress has been made in developing ML parametrizations which are stable within realistic
152 models (i.e. not toy models, aquaplanets etc.), and ML-based full models which can run stably and skillfully to longer
153 lead times. This is usually achieved through training the model on more comprehensive data, employing ML
154 architectures which keep the model outputs within physically real limits, or imposing physical constraints or
155 conservation rules within the ML architecture or training loss functions^f.

156 There are still challenges and possible limitations to an ML approach to weather and climate modelling. In most cases,
157 a robust ML model or parameterization scheme should be able to:

- 158 • remain stable in a full (i.e. non-idealized) model run,
- 159 • generalize to cases outside its training envelope,
- 160 • conserve energy and achieve the required closures.

161 Additionally, for an ML approach to be worthwhile it must provide one or more of the following benefits:

- 162 • For ML parametrization schemes:
 - 163 ○ a speedup of the representation of a ~~subgrid-scale~~ process vs. when run with a traditional
164 parametrization scheme. This can make the difference between the scheme being cost-effective to
165 run or not - when it is not cost-effective the process usually needs to be represented with a static
166 forcing or boundary condition file,
 - 167 ○ a speedup of the model vs. when run with traditional parametrization schemes,
 - 168 ○ improved representation of sub-grid process(es) over traditional parameterization schemes, as
169 measured by metrics appropriate to the situation,
 - 170 ○ improved overall accuracy/skill of the model when run with traditional parametrization schemes,
 - 171 ○ insight into physical processes not provided by current numerical models or theory.
- 172 • For full ML models:
 - 173 ○ a speedup of the model vs. an appropriate numerical model control,
 - 174 ○ improved overall accuracy/skill of the model vs. an appropriate numerical model control,
 - 175 ○ skillful prediction to greater lead times than an appropriate numerical model control,

Deleted: because

Deleted: can

Formatted: Superscript

Deleted: sub-grid scale

179 o insight into physical processes not provided by current numerical models or theory
180 Furthermore, in some cases of ML approaches to weather and climate modelling problems (particularly for full model
181 replacement) the work is led by data scientists and ML researchers with limited expertise in weather and climate model
182 evaluation. This can lead to flawed, misleading or incomplete evaluations. Hewamalage et al. (2022) have sought to
183 rectify this problem by providing a guide to forecast evaluation for data scientists.

Deleted: many

184 The scope of this review is deliberately limited to the application of ML within numerical weather and climate models
185 or for their replacement. This is done to keep the length of this review manageable. ML has enormous utility for other
186 aspects of the forecast value chain such as observation quality assurance, data assimilation, model output
187 postprocessing, forecast/product generation, downscaling, impact prediction, decision support tools, etc. A review of
188 the application of, and progress in, ML in these areas would be of great value but is outside the scope of this review
189 and is left to other work. Molina et al. (2023) have provided a very useful review of ML for climate variability and
190 extremes which is highly complementary to this review. They draw similar lines of delineation in the earth system
191 modelling (ESM) value chain to those mentioned above; describing them as “initializing the ESM, running the ESM,
192 and postprocessing ESM output”. They examine each of these steps in turn, with a focus on the prediction of climate
193 variability and extremes. Here we take a different approach, focusing on one part of the value chain (running the
194 ESM), but looking in more detail at this one part. Additionally, here we consider climate modelling in the context of
195 multiyear and free-running multidecadal simulations, but exclude the topic of ML for climate change projections,
196 climate scenarios, and multi-sector dynamics. This is again in the interests of ensuring the scope of the review is
197 manageable, rather than because these topics are not worthy of review. On the contrary, a review dedicated to the
198 utility of machine learning in this area would be of enormous value to the community, but cannot be adequately
199 explored here. A brief introduction to key ML architectures and concepts, including suggested foundational reading,
200 is also provided to aid readers who are unfamiliar with the subject.

Deleted: for future

Deleted:

201 The remainder of this review is structured as follows: In Section 2 an introduction to the two ML architectures most
202 prevalent in the review is provided, followed by a suggested methodological approach to applying ML to a problem,
203 and finally a brief overview of some of the major ML architectures and algorithms. With this background in place, the
204 application of ML in weather and climate modelling is explored in the following five sections. In Section 3, ML use
205 in sub-grid parametrization and emulation, along with tools and challenges specific to this domain, are covered.
206 Zooming out from subgrid-scale to processes resolved on the model grid, in Section 4 the application of ML for the
207 partial differential equations governing fluid flow is reviewed. Expanding scope further, again to consider the entire
208 system, the use of ML for full model replacement or emulation is reviewed in Section 5. In Section 6 the growing field
209 of physics constrained ML models is introduced, and in Section 7 a number of topics tangential to the main focus of
210 this review are briefly mentioned. Setting the work covered in the previous sections in a broader context, a review of
211 the history of, and progress in, ML outside of the fields of weather and climate science is presented in Section 8. In
212 Section 9 some practical considerations for the integration of ML innovations into operational and climate models are
213 discussed, followed by a short introduction to some of the ethical considerations associated with the use of ML in
214 weather and climate modelling in Section 10. In Section 11, some future research directions are speculated on, and
215 some suggestions are made for promising areas for progression. Finally, a summary is presented in Section 12, and a

Deleted: could

Deleted: not

Deleted: done proper justice

Deleted: to

Deleted: quick

Deleted: and methodologies is provided

Deleted: ,

Deleted: before

Deleted: .

Deleted: sub-grid scale

Deleted: yet

Deleted: and f

Deleted: 0

Deleted: . A

233 Glossary of Terms is provided after the final Section to aid the reader in their understanding of key concepts and
234 words.

235

236 2. A Quick Introduction to Machine Learning

237 While the scope of this paper is a review of ML work directly applicable to weather and climate modelling, an abridged
238 introduction to some key fundamental ML concepts is provided here to aid the reader. Suggested starting points for
239 interested readers, including guidance on the utility of different model architectures and algorithms, and the
240 connections between different applications and approaches, are as follows:

- 241 • Hsieh (2023) provides a thorough textbook on environmental data science including statistics and machine
242 learning
- 243 • Chase et al (2022a, 2022b) provide an introduction to various machine learning algorithms with worked
244 examples in a tutorial format and an excellent on-ramp to ML for weather and climate modelling.
- 245 • Russell & Norvig (2021) provide a comprehensive book regarding artificial intelligence in general.
- 246 • Goodfellow et al. (2016) provide a well-regarded book on deep learning theory and modern practise.
- 247 • Hastie et al. (2009) provide a book on statistics and machine learning theory.

248 This introductory section is a brief exposition of the concepts most central to this review. Definitions for this section
249 can be found in the glossary.

250 The majority of ML methods which have found traction in weather and climate modelling were first developed in
251 fields such as computer vision, natural language processing and statistical modelling. Few, if any, of the methods
252 mentioned in this paper could be considered unique to weather and climate modelling, however, they have in many
253 cases been modified to a greater or lesser extent to suit the characteristics of the problem. In this review, the term
254 algorithm refers to the mathematical underpinnings of a machine learning approach. By this definition, decision trees
255 (DTs), NNs, linear regression and Fourier transforms are examples of algorithms. The two most relevant algorithms
256 for this review are DTs and NNs. Many ML algorithms can be thought of as optimizing a nonlinear regression, with
257 deep learning utilizing an extremely high-dimensional model. There is no consensus on the definition of ML, with the
258 term encompassing relatively large or small topical domains depending on who is asked. A good rule of thumb,
259 however, is that any iterative computational process that seeks to minimize a loss function or optimize an objective
260 function can be considered to be a form of ML. Some of the chief concerns in machine learning are generalizability
261 of the models, how to train (optimise the variables of) the model, and how to ensure robustness. The inputs and outputs
262 of machine learning models are the often same as physical models or model components. The term architecture in
263 machine learning refers to a specific way of utilizing an algorithm to achieve a modelling objective reliably. For
264 example, the U-Net[†] architecture is a specific way of laying out a NN which has proven effective in many applications.
265 The extreme gradient boosting decision tree[‡] architecture is a specific way of utilizing DTs which has proven reliable
266 and effective for an extraordinary number of problems and situations and is an excellent choice as a first tool to
267 experiment with machine learning.

Deleted: ¶

Deleted: The focus of the review is not specifically on architecture selection. This is in part because of the lack of a clear finding associated with such an analysis, and also because the taxonomy and classification of model types is not necessarily straightforward. The chart below shows that by far the two most frequently used general categories of architecture are fully connected NNs and convolutional neural networks of various sub-types. Some of the most significant research findings come from entirely novel architectures which by definition cannot have wide adoption yet. ¶

[1]

Deleted: s¶

Deleted: include

Deleted: ,

Deleted: ,

Deleted:

Deleted: , and

Deleted: .

Deleted: Furthermore, there is a trend towards increasingly customized architectures in this field as it matures. ¶

Deleted: machine learning

Deleted:

Deleted: one-size-fits all simple mental model for all machine learning

Deleted: .

Deleted: set

Deleted: ¶

311 ~~A major current focus of ML research in the context of weather and climate modelling is new NN-based architectures~~
312 ~~and algorithms, and improved training regimes.~~ Many other algorithms have been and continue to be employed in
313 machine learning ~~more broadly,~~ but are not ~~pertinent~~ to this review.

Deleted: A major current focus of ML research is new architectures based on NNs. Research also continues into the algorithms themselves.

Deleted: central

314 A key point for ML researchers to be aware of is the critical importance of approaching model training carefully.
315 There are many pitfalls which can result in underperformance, unexpected bias or misclassification. For instance,
316 adversarial examples[†] can occur ‘naturally’, and systems which process data can be subject to adversarial attack[†]
317 through the intentional supply of data designed to fool a trained network.

318 2.1. Introduction to Neural Networks

319 NNs can be regarded as universal function approximators (Hornik et al., 1989; see also Lu et al., 2019). Further, NN
320 architectures can theoretically be themselves modelled as a very wide feed-forward[†] NN with a single hidden layer.
321 A Fourier transform is another example of a function approximator, although it is not universal since not all functions
322 are periodic. ~~NNs can therefore theoretically be candidates for accurate modelling of physical processes, although in~~
323 ~~practise they cannot always reliably interpolate beyond their training envelope and as such may not generalize to new~~
324 ~~regimes.~~ ML models are typically introduced in the literature as being either classification[†] or regression[†] models, and
325 either supervised[†] or unsupervised[†].

Deleted: NNs can therefore be candidates for accurate modelling of physical processes.

Deleted: ¶

326 The mathematical underpinning of a NN can be considered distinctly in terms of its evaluation[†] (i.e., output, or
327 prediction) step and its training update step. The prediction step can be considered as the evaluation of a many-
328 dimensional arbitrarily complex function.

329 The simplest NN is a single-input, single node network with a simple activation[†] function. A commonly used activation
330 function ~~for a single neuron~~ is the sigmoid function, which helpfully compresses the range between 0 and 1 ~~while~~
331 ~~allowing a nonlinear response.~~ A classification model will employ a threshold to map the output into the target
332 categories. A regression model seeks to optimize the output result against some target value for the function. ~~Larger~~
333 ~~networks make more use of linear activations and may utilise heterogenous activation function choices at different~~
334 ~~layers.~~

335 Complex NNs are built up from many individual nodes, which may have heterogenous activation functions and a
336 complex connectome[†]. The forward pass[†], by which inputs are fed into the network and evaluated against activation
337 functions to produce the final prediction, uses computationally efficient processes to quickly produce the result.

338 The training step for a NN is far more complex. The earliest NNs were designed by hand rather than through
339 automation. The training step applies a back-propagation[†] algorithm to apply adjustment factors to the weights[†] and
340 biases[†] of each node based on the accuracy of the overall prediction from the network.

341 Training very large networks was initially impractical. Both hardware and architecture advances have changed this,
342 resulting in the significant increase in application of NNs to practical problems. Most NN research explores how to
343 utilize different architectures to train more effective networks. There is little research going into improving the
344 prediction step as the effectiveness of a network is limited by its ability to learn rather than its ability to predict. Some

352 research into computational efficiency is relevant to the predictive step. NNs can still be technically challenging to
353 work with, and a lot of skill and knowledge are needed to approach new applications.

354 The major classes of NN architectures most likely to be encountered are:

- 355 • ~~Small, fully-connected networks, which are less commonly featured in recent publications but are still~~
356 ~~effective for many tasks and are still being applied and may well be encountered in practice.~~
- 357 • Convolutional[†] architectures, first applied to image content recognition, which match the connectome of the
358 network to the fine structure of images in hierarchical fashion to learn to recognize high-level objects in
359 images
- 360 • ~~Recurrent token-sequence architectures, first applied to natural language processing, generation and~~
361 ~~translation; applicable to any time-series problem. Now also applied to image and video applications, and~~
362 ~~mixed-mode applications such as text-to-image or text-to-video~~
- 363 • ~~Transformer architectures[†], based on the attention mechanism[†] to provide a non-recurrent architecture which~~
364 can be trained using parallelized training strategies. This allows larger models to be trained. Originally
365 developed for sequence prediction and extended to image processed through vision transformer architectures.

Deleted: <#>Token-sequence architectures, first applied to natural language processing, generation and translation; applicable to any time-series problem. Now also applied to image and video applications, and mixed-mode applications such as text-to-image or text-to-video¶

Deleted: <#>.

Deleted: <#>¶

366 2.2. Introduction to Decision Trees

367 DTs are a series of decision points, typically represented in binary fashion based on a simple threshold. A particular
368 DT of a particular size maps the input conditions into a final 'leaf' node which represents the outcome of the decisions
369 up to that point.

370 A random forest[†] (RF) is the composition of a large number of DTs assembled according to a prescribed generation
371 scheme, which are used as an ensemble. A gradient boosted decision tree (GBDT) is built up sequentially, where each
372 subsequent decision tree attempts to model the errors of the stack of trees built up thus far. This approach outperforms
373 RFs in most cases.

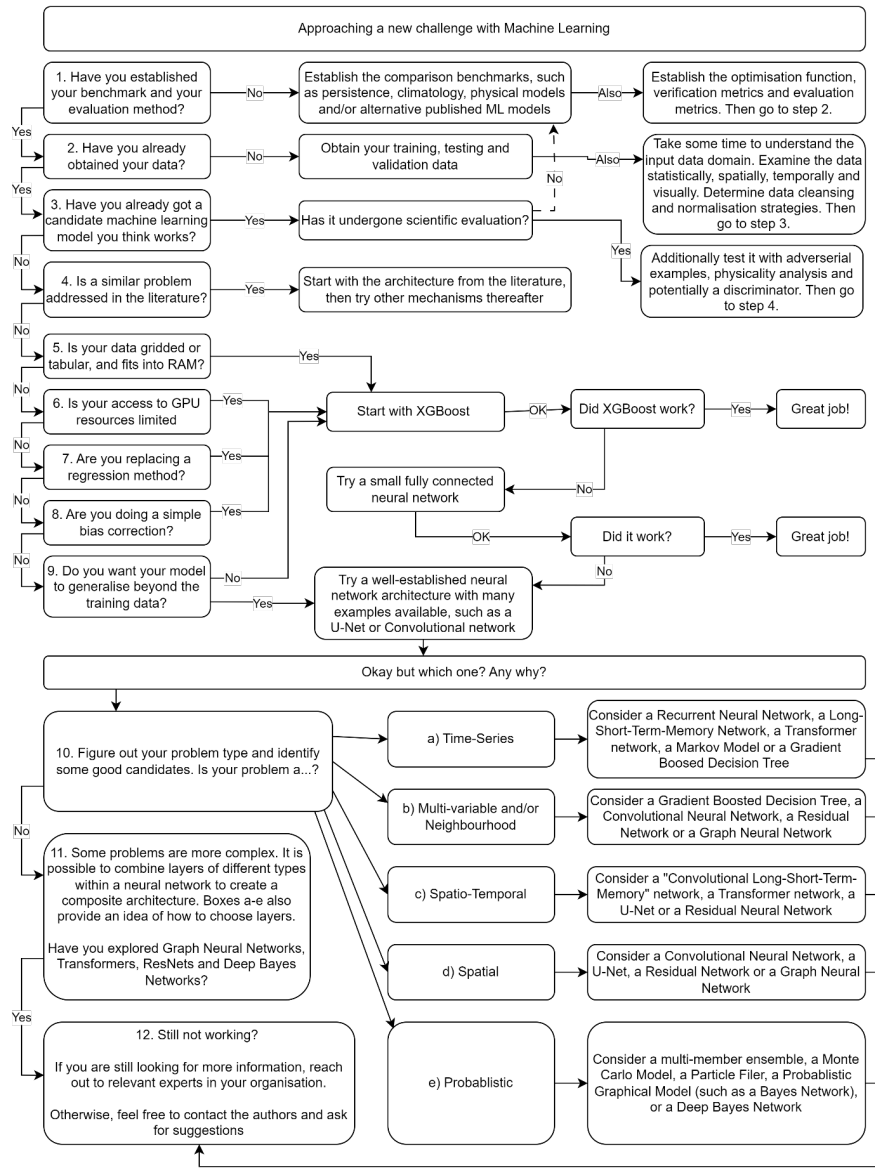
374 The DT family of ML architectures are very easy to train and are very efficient. They are well documented in the
375 public domain and in published literature. DTs are statistical in nature and are not capable of effectively generalizing
376 to situations which are not similar to those seen during training. This can be an advantage when unbounded outputs
377 would be problematic, however can lead to problems where an ability to produce out-of-training solutions is necessary.
378 Additionally, current DT implementations require all nodes (of all trees in the case of RFs and GBDTs) to be held in
379 memory at inference time, making them potentially memory heavy.

380 2.3. Methodologies for Machine Learning

381 ~~It is challenging to provide simplified advice for how to approach problem-solving in ML. There are few strict~~
382 ~~theoretical reasons to choose any one of the variety of architectures which are available. The authors would also~~
383 ~~caution against assuming that results in the literature are the product of a detailed comparison of alternative~~
384

Deleted: machine learning

Deleted: result



394
395

Figure 2: A methodological flowchart illustrating a suggested approach to applying ML to a research problem.

396 architectures, or assuming that a deep learning approach is going to be easy or straightforward. It will often be the
 397 case that multiple machine learning architectures may be similarly effective, and determining the optimal
 398 architecture is likely to involve extensive iteration. Any specific methodology is also likely to reflect the intuitions
 399 (or biases), knowledge, and background of the authors of that methodology.
 400 Nonetheless, there is an appetite from many scientists for reasonable ways to 'get started' and to provide some
 401 assistance for practical decision-making, particularly if approaching the utilization of machine learning for the first
 402 time or in a new way. Figure 2 provides a set of suggested steps and decision points to help readers approach a new
 403 challenge with ML.
 404 The flowchart presented in Figure 2 provides an overview of methodological steps that can be taken when using ML
 405 to solve a problem, however it does not give much insight into the pros and cons of the common ML architectures
 406 available and used in the literature. Table 1 provides a brief summary of the major ML architectures and algorithms
 407 used by the studies cited in this review and gives a short note on some of their pros and cons. This table is not
 408 exhaustive, and readers are strongly encouraged to use it as a starting point for further exploration, rather than a
 409 definitive guide. The relative strengths and weakness of each ML architecture can be subtle, and highly dependent on
 410 the use case, their application, and their tuning. Establishing a good understanding of the ML architecture being used
 411 is a critical step for any scientist intending to delve into ML modelling. Interested readers should also refer to Chase
 412 et al (2022b), where a similar table is presented that covers a wider variety of traditional methods but fewer neural
 413 network approaches.
 414 An increasingly diverse array of ML architectures are being applied to an ever-growing variety of challenges. These
 415 architectures all have sub-variants and ancestor architectures which may not be represented, all of which may be found
 416 to be of use for weather and climate modelling applications. Other concerns, such as data normalization, training
 417 strategies, and capturing physicality become as relevant as the choice of architecture once a certain level of
 418 performance is achieved.
 419 Figure 3 shows a summary of the ML architectures and algorithms used by the studies cited in this review, including
 420 the number of times each architecture is used. It can be seen from this that the two most frequently used general
 421 categories of architecture are Fully Connected NNs (FCNNs) and Convolutional NNs (CNNs) of various sub-types.
 422 However, some of the most significant recent research findings come from new architectures which by definition
 423 cannot have wide adoption yet (these are grouped under the 'Mixed/custom NN' category in Figure 3).
 424 In some cases, little justification is given for the ML architecture used in a study, and readers are therefore cautioned
 425 against using the relative popularity of a particular ML architecture in the literature as a guide for its suitability for a
 426 given task.
 427 Furthermore, ML models increasingly use a mix of different algorithms and architectures. For example, a common
 428 combination is fully-connected NN layers, convolutional NN layers, and LSTM layers. For the purposes of Figure 3,
 429 the authors have endeavoured to categorise the ML architectures used in the studies in this review as accurately as
 430 possible, with complex architectures being placed in the "Mixed/custom NN" category, however, where an
 431 architecture was mostly but not entirely aligned with a single category, it was placed in that category. For example,
 432
 433

Deleted:

Deleted: ¶
Comcomittantly, a

Deleted: aka

Formatted: Line spacing: 1.5 lines

Deleted: ¶
The authors therefore suggest the following approach may be of assistance:

Deleted: ... [2]

Deleted: Further to the

Deleted: above

Deleted: ,

Deleted: more guidance is provided to explore the architectures used in papers contained in this review an ... [3]

Deleted: ¶
Approach ... [4]

Deleted: ¶

Deleted: ¶

Deleted: ever-increasing variety

Deleted: of neural network

Deleted: increasing

Deleted: turn out to be the next step of the journey

Moved (insertion) [2]

<u>Approach</u>	<u>Description</u>	<u>Pros</u>	<u>Cons</u>
<u>Simple regression techniques</u>	<u>Includes linear regression and logistic regression. See Chase et al. (2022b) for more detail.</u>	<u>Explainable and well-understood.</u>	<u>Can only capture simple relationships.</u>
<u>Decision Tree</u>	<u>Consists of a series of branching decisions, culminating in a number of decision 'leaves'. The decision points are trainable.</u> <u>Provides the basis for understanding more complex decision tree and regression tree approaches.</u>	<u>Easily explainable.</u> <u>Computationally tractable and fast</u>	<u>Unable to fully model complex problems.</u> <u>Cannot make predictions outside the training envelope.</u>
<u>Random Forest (RF)</u>	<u>A random forest consists of many decision trees, which form an ensemble and the average result is taken. The construction of the trees uses randomness.</u>	<u>Versatile and effective.</u> <u>Computationally tractable and fast.</u> <u>Allows focus on the input variables rather than on process or model definition.</u>	<u>Usually performs slightly less well than gradient boosted decision trees.</u>
<u>Gradient Boosted Decision Trees (GBDT)</u>	<u>Akin to Random Forests, however each additional member is used to predict the residual error of the ensemble so far.</u> <u>Is often sufficient for a given problem, and should thus be considered as a baseline for measuring more complex ML models against.</u>	<u>A highly versatile and reliable approach.</u> <u>Computationally tractable and fast.</u> <u>Allows focus on the input variables rather than on process or model definition.</u> <u>Feature importance plots can guide intuition.</u>	<u>Has practical limitations at scale due to large memory requirements at inference time.</u> <u>Limited ability to simulate complex systems compared to other ML approaches such as NNs.</u> <u>Cannot make predictions outside the training envelope without customized leaves.</u>

<u>Vector Machines</u>	<u>Support Vector Machines (SVMs) and Relevance Vector Machines (RVMs) are supervised models used for regression and classification. RVMs have the same functional form as SVMs, but are a probabilistic classification based on Bayesian inference. Vector Machines seek to define the optimal division between classes by finding the hyperplanes which have the largest distance to the nearest training-data point of any class.</u>	<u>Can be used for similar problems as GBDTs. Computationally efficient and often effective. Mathematically appealing. Capable of modelling nonlinear functions.</u>	<u>Now less-used compared to random forests and GBDTs.</u>
<u>Single neuron</u>	<u>See Chase et al. (2022b) for a description of the structure of a perceptron. Forms the conceptual and structural basis for all NN architectures.</u>	<u>Unused in practice outside of a larger NN architecture.</u>	<u>Unable to model most problems in isolation.</u>
<u>Fully-Connected feed-forward Neural Network (FCNN)</u>	<u>Consists of multiple layers of neurons, with each neuron being connected to every neuron in the subsequent layer. Still quite widely used in weather and climate modelling, in spite of declining use in other machine learning domains. Is often sufficient and should be considered as a baseline for measuring more complex architectures against.</u>	<u>Effective for applications such as parametrization scheme emulation and PDE solver preconditioning. Relatively simple to work with. Computationally tractable.</u>	<u>Unable to effectively train beyond a certain size or depth, and thus is increasingly being replaced with more complex architectures as ML moves to deeper NNs.</u>

<u>Bayesian networks</u>	<u>A system (probabilistic graphical model) comprised of nodes which together predict both an expected value and a likelihood. Each node is associated with a probability function that provides a probability (or distribution) of the variable represented by the node.</u>	<u>Effective for refining an expert or knowledge-based model by incorporating additional observations. Capable of dealing with both semantic concepts and physical processes.</u>	<u>Determining an optimal model can be challenging and training times are prohibitive for large networks.</u>
<u>Deep Bayesian Networks</u>	<u>Deep Bayesian techniques attempt to capture the model complexity of deep neural networks while retaining the ability to predict a distribution of outcomes, a probabilistic model and a clear information-theoretical bases.</u>	<u>Used to obtain a more realistic expression of uncertainty. Effective in modelling where causal relationships aren't understood.</u>	<u>Not as well explored as neural networks in recent literature.</u>
<u>Convolutional Neural Network (CNN)</u>	<u>Involves convolving a (usually 2D image, but can also be 1D temporal, for example) input field with a filter function (often a top hat function') to extract features on different spatial scales. Conceptually useful in understanding how a neural network can build up an abstract or 'big picture' definition of a process in its hidden layers by assembling fine-scale features.</u>	<u>The go-to network for image-based problems. Proven effective on many problems and is well-covered in the literature.</u>	<u>May require more significant hardware such as a modern GPU.</u>
<u>Residual Neural Network (ResNet)</u>	<u>ResNets are a form of CNN including skip connections, whereby the inputs of a number of convolutional layers are appended to the outputs of those layers to retain information lost through the weights in the convolutional layers. These skip connections make it possible to train much deeper convolutional networks than would be possible otherwise.</u>	<u>Allows very deep networks to be efficiently trained. Allows an iterative build-up of network size by experimenting with the number of residual layers. Could be a good choice to couple with physically interpretable layers.</u>	<u>Somewhat more computationally costly than other deep architectures.</u>

<p><u>U-Net</u></p>	<p><u>Derives its name from the shape of the network as it is commonly shown diagrammatically (it forms a “U” shape).</u></p> <p><u>Consists of a series of downsampling convolutional layers, each of which further abstracts the information in the inputs (forming the first half of the “U”). These are then upsampled again to the original resolution of the input data (forming the second half of the “U”). Each downsampling step has its output appended to the input of the corresponding upsampling step (a form of skip connection).</u></p>	<p><u>Effective for many purposes and widely used in classification and image segmentation. Has also seen uptake for nowcasting applications and prediction of multiyear timescale ocean variables.</u></p>	<p><u>No serious drawbacks. Has somewhat given way to more complex architectures recently</u></p>
<p><u>Deep Operator Network (DeepONet)</u></p>	<p><u>A NN which is designed to learn the mappings between inputs and outputs of the mathematical operators underpinning processes, rather than directly predicting the outputs of the processes themselves. Was developed in the context of fluid dynamics and differential operators.</u></p> <p><u>An important theoretical component of the Adaptive Fourier Neural Operator used in FourCastNet (Pathak et al., 2022).</u></p>	<p><u>Provides a strong theoretical basis for learning the underlying function space of a data set.</u></p> <p><u>Highly effective for fluid dynamics and idealized systems.</u></p> <p><u>Can retain the properties of the learned operators. For example, can exhibit translational and scale invariance where that property holds for the operator in question.</u></p>	<p><u>Conceptually not straightforward.</u></p> <p><u>Requires strong mathematical and machine learning expertise to apply effectively to new challenges.</u></p>

<u>Graph Neural Network (GNN)</u>	Models data as a set of interconnected nodes and edges (as opposed to assuming data is on a regular grid). Underpins Keisler (2022) and GraphCast (Lam et al., 2022)	Does not require data to be on a grid or distributed in a uniform manner. Capable of incorporating teleconnections, nonlocal relationships, and other complex variable relationships.	Costly to train.
<u>Discriminator</u>	A NN is trained to discriminate between two examples and identify the “real” one. Is used to estimate whether a sample is from the observations or the model. Forms one part of a GAN.	Can be used in place of a manually-defined loss function to train without over-emphasizing any individual metrics or variables. Can be used as an effective loss function when training Can be used independently to evaluate model realism. Comes closest to human subjective evaluation of image quality.	Is more likely to require more machine learning domain knowledge to resolve issues.
<u>Generative Adversarial Network (GAN)</u>	Combines a generator network with a discriminator and trains them in an adversarial manner: the discriminator tries to differentiate the generator from ground truth, the generator tries to trick the discriminator. Eventually the discriminator can't differentiate the generator from ground truth. May be part of a multi-phase training strategy in order to improve realism after initial optimization.	Produce results which prioritize realism over accuracy (could also be a con). Is less prone to the blurring that results from training to simpler loss functions and thus can be more effective in producing sharp images and predicting statistical extremes.	Increases training costs. Favors a ‘good looking’ answer over a correct answer. Can be difficult to train as the generator and discriminator must be kept balanced (one can outperform the other leading to mode collapse – a false minima).

<p><u>Recurrent Neural Network (RNN)</u></p>	<p><u>Any neural network where the output of previous predictions are provided to a sequence-based model. Multiple sub-types of the RNN exist.</u></p>	<p><u>A simple RNN design can model many problems effectively.</u> <u>A recurrent architecture allows access to and inspection of the belief state at each iteration.</u></p>	<p><u>Recurrent approaches can accumulate errors quickly.</u> <u>Relationships which act over longer time-frames or distances than the recurrence length may not be captured.</u> <u>Choosing the length of the sequence may be a challenge.</u></p>
<p><u>Long Short Term Memory (LSTM) Network</u></p>	<p><u>Contains modified neurons with a memory component and the ability to retain or forget information. Is applied to sequence inputs and can learn the sequential scales in which information is encoded (e.g., what timescales in a timeseries are pertinent for future prediction).</u> <u>Has been combined with the ideas underpinning CNNs to create Convolutional LSTMs (ConvLSTM), which fit for both timescales of relevance and spatial features of relevance.</u></p>	<p><u>An effective alternative to a recurrent network which has proven very good at modelling time-series.</u> <u>A proven and effective mechanism for dimensionality reduction to allow the training of large networks.</u></p>	<p><u>May not include spatial relationships (unless it's a ConvLSTM), and may be more complex than needed for some problems.</u> <u>Less explainable than an attention mechanism.</u> <u>Has a bias towards closer points in a sequence (e.g., will be biased towards the recent past over a longer timescale in time series prediction).</u></p>

<p><u>Attention Mechanism</u></p>	<p><u>Often used in conjunction with other architectures as a feature extraction/dimensionality reduction method.</u></p> <p><u>A NN is trained to learn the degree of importance of each input datapoint on each other one in a sequence.</u></p> <p><u>Attention mechanism-based NNs are rapidly overtaking LSTMs as the method of choice for modelling sequence-based information.</u></p>	<p><u>Unlike LSTMs, attention mechanisms are not biased towards relationships between near points in a sequence.</u></p> <p><u>Rather, attention mechanisms treat all points in an input sequence equally and retain the learned attention mappings between each point.</u></p> <p><u>In the context of weather and climate modelling, the learned attention mappings between points can be a useful tool for assessing the degree to which a NN has learned physically realistic teleconnections.</u></p>	<p><u>More costly to train than an LSTM for the same problem because attention mechanisms have more free parameters.</u></p>
<p><u>Transformer</u></p>	<p><u>The transformer architecture combines an attention mechanism with an autoregressive approach whereby each previously predicted step in a sequence is an input into the prediction of the next step.</u></p> <p><u>Transformer architectures underpin the current generation of language models such as ChatGPT.</u></p> <p><u>Transformers are now often included as part of other architectures for input dimensionality reduction.</u></p>	<p><u>A proven and effective mechanism for dimensionality reduction to allow the training of large networks.</u></p> <p><u>While the uptake of transformer architectures in weather and climate modelling is still small, their impressive performance for sequence prediction suggests they could have great for the field.</u></p>	<p><u>Transformers can be difficult to train due to a tendency to overemphasize the recurrent component of the network over new inputs in the early stages of training.</u></p>

Table 1: A summary of major ML architectures and algorithms used by the studies cited in this review. Interested readers should also refer Chase et al (2022b) where a similar table is presented that covers a wider variety of traditional methods but fewer neural network approaches.

an LSTM model with a small number of feed-forward layers would be categorised as a Recurrent NN. Since many contemporary ML models combine multiple architectural elements and algorithms into the one model, it is somewhat of an oversimplification to consider each of these in isolation, and while starting with a simple model design with a

Formatted: Font: English (US)

Moved up [2]: Furthermore, ML models increasingly use a mix of different algorithms and architectures. For example, a common combination is fully-connected NN layers, convolutional NN layers, and LSTM layers. For the purposes of Figure 3, the authors have endeavoured to categorise the ML architectures used in the studies in this review as accurately as possible, with complex architectures being placed in the "Mixed/custom NN" category, however, where an architecture was mostly but not entirely aligned with a single category, it was placed in that category. For example,

470
471
472
473
474
475
476
477


488 limited selection of layer types is advisable to aid interpretability, there is no reason they cannot be combined or used
 489 in conjunction with each other if this improves the performance of the model.
 490 Adapting, optimizing and debugging issues with machine learning systems can be very complex (especially so for
 491 large NNs), and is likely to require both machine learning expertise and domain knowledge (i.e. scientific knowledge).
 492 XGBoost provides the ability to generate chart showing the importance of the features in the model which can be very
 493 helpful. Shapley Additive Explanations (Lundberg and Lee 2017) can provide insights into feature importance for any
 494 model including NNs.

Formatted: Font: (Default) Times New Roman, 10 pt, Font colour: Auto, English (AUS)

Commented [CC2]: This last para feels a bit out of place...

Deleted: <object>

Formatted: Indent: Left: 0 cm



496
 497 Figure 3: A count of the ML architectures and algorithms used by the studies cited in this review. As with Figure 1, this figure
 498 includes all references from this review except for: seminal ML papers that are on new ML methods (e.g., foundational ML papers),
 499 review papers, any paper cited that concerns a topic which is out of scope (e.g., nowcasting), and any other paper which does not
 500 present a new method directly applicable to weather and climate modelling. The full table of citations is provided in the appendix,
 501 ↓

Formatted: Font: 9 pt, English (AUS)

Deleted: ¶
 Adapting, optimizing and debugging issues with machine learning systems can be very complex (especially so for large NNs), and is likely to require both machine learning expertise and domain knowledge (i.e. scientific knowledge). XGBoost provides the ability to generate chart showing the importance of the features in the model which can be very helpful. Shapley Additive Explanations (Lundberg and Lee 2017) can provide insights into feature importance for any model including NNs.

Deleted: Sub-grid scale

Deleted: sub-grid scale

Deleted: sub-grid scale

502 **3. Sub-grid parametrization and emulation**

503 Subgrid-scale processes in numerical weather and climate models are typically represented via a statistical
 504 parameterization of what the macroscopic impacts of the process would be on resolved processes and parameters.
 505 These are commonly referred to as parameterization schemes, and can be very complex and relatively computationally
 506 costly. For example, in the European Centre for Medium-Range Weather Forecast's (ECMWF) Integrated Forecasting
 507 System (IFS) model they account for about a third of the total computational cost of running the model (Chantry et al.
 508 2021b). They also require some understanding of the underlying unresolved physical processes. Examples of subgrid-
 509 scale processes which are typically currently parameterized in operational systems include gravity wave drag,
 510 convection, radiation, subgrid-scale turbulence, and cloud microphysics. As additional complexity (for example

525 representation of aerosols, atmospheric chemistry, land surface processes, etc.) is added to numerical models, the
526 computational cost will only increase.

527 ML presents an alternative approach to representing subgrid-scale processes, either by emulating the behavior of an
528 existing parametrization scheme, emulating the behavior of sub-components of the scheme, by replacing the current
529 scheme or sub-component entirely with an ML-based scheme, or by replacing the aggregate effects of multiple
530 parametrization schemes with a single ML model.

531 ML emulation of existing schemes or sub-components has the advantage of maintaining the status quo within the
532 model; no or minimal re-tuning of the model should be required since the ML emulation is trained to replicate the
533 results of an already-tuned-for scheme. Because of this, the main benefit of this approach is that it reduces the
534 computational cost of running the parametrization scheme. On the other hand, full replacement of an existing
535 parameterization scheme or sub-component with an ML alternative has the potential to be both computationally
536 cheaper and also an improvement over the preceding scheme.

537 In the following subsections, a review of the literature on aspects of ML for the parametrization and emulation of
538 subgrid-scale processes is presented.

Deleted: sub-grid scale

Deleted: sub grid-scale

539 3.1. Early work on ML parametrization and ML emulations

540 A popular target for applying ML in climate models is radiative transfer, since it is one of the more computationally
541 costly components of the model. As such, many early examples of the use of ML in sub-grid parametrization schemes
542 focus on aspects of this physical process. Chevallier et al. (1998) trained NNs to represent the radiative transfer budget
543 from the top of the atmosphere to the land surface, with a focus on application in climate studies. They incorporated
544 the information from both line-by-line and band models in their training to achieve competitive results against both
545 benchmarks. Their NNs achieved accuracies comparable to or better than benchmark radiative transfer models of the
546 time, while also being much faster computationally.

547 In contrast to the ML based scheme developed by Chevallier et al. (1998), which could be considered an entirely new
548 parametrization scheme, Krasnopolsky et al. (2005) used NNs to develop an ML based emulation of the existing
549 atmospheric longwave radiation parametrization scheme in the NCAR Community Atmospheric Model (CAM). The
550 authors demonstrated speedups with the NN emulation of 50-80 times the original parameterization scheme.

551 Emulation of existing schemes has since then become a popular method for achieving significant model speedups. For
552 example, Gettelman et al. (2021) investigated the differences between a General Circulation Model (GCM) with the
553 warm rain formation process replaced with a bin microphysical model (resulting in a 400% slowdown) and one with
554 the standard bulk microphysics parameterization in place. They then replaced the bin microphysical model with a set
555 of NNs designed to emulate the differences observed, and showed that this configuration was able to closely reproduce
556 the effects of including the bin microphysical model, without any of the corresponding slowdown in the GCM.

559 **3.2. ML for coarse graining**

560 Coarse graining involves using higher resolution model or analysis data to map the relationship between smaller scale
561 processes and a coarser grid resolution. It can be used to develop parameterization schemes without explicitly
562 representing the physics of smaller scale processes.

563 This has proven to be a popular method for developing ML-based parametrization schemes. Brenowitz & Bretherton
564 (2018) used a near-global aqua planet simulation run at 4 km grid length to train a NN to represent the apparent sources
565 of heat and moisture averaged onto 160 km² grid boxes. They then tested this scheme in a prognostic single column
566 model and showed that it performed better than a traditional model in matching the behavior of the aqua planet
567 simulation it was trained on. Brenowitz & Bretherton (2019) built on this work by training their NN on the same global
568 aqua-planet 4 km simulation, but then embedded this scheme within a coarser resolution (160 km²) global aqua planet
569 GCM. Embedding NNs within GCMs is challenging because feedbacks between NN and GCM components can cause
570 spatially extended simulations to become dynamically unstable within a few model days. This is due to the inherently
571 chaotic nature of the atmosphere in the GCM responding to inputs from the NN which cause rapidly escalating
572 dynamical instabilities and/or violate physical conservation laws. The authors overcame this by identifying and
573 removing inputs into the NN which were contributing to feedbacks between the NN and GCM (Brenowitz et al. 2020),
574 and by including multiple time steps in the NN training cost function. This resulted in stable simulations which
575 predicted the future state more accurately than the course resolution GCM without any parametrization of subgrid-
576 scale variability, however the authors do observe that the mean state of their NN-coupled GCM would drift, making
577 it unsuitable for prognostic climate simulations.

578 Rasp et al. (2018) trained a deep NN to represent all atmospheric subgrid processes in an aquaplanet climate model
579 by learning from a multiscale model in which convection was treated explicitly. They then replaced all sub-grid
580 parameterizations in an aquaplanet GCM with the deep NN, and allowed it to freely interact with the resolved
581 dynamics and the surface-flux scheme. They showed that the resulting system was stable and able to closely reproduce
582 not only the mean climate of the cloud-resolving simulation but also key aspects of variability in prognostic multiyear
583 simulations. The authors noted that their decision to use deep NNs was a deliberate one, because they proved more
584 stable in their prognostic simulations than shallower NNs, and they also observed that larger networks achieved lower
585 training losses. However, while Rasp et al. (2018) were able to engineer a stable model that produced results close to
586 the reference GCM, small changes in the training dataset or input and output vectors quickly led to the NN producing
587 increasingly unrealistic outputs and causing model blow-ups (Rasp 2020). Consistent with this, Brenowitz &
588 Bretherton (2019) report that they were unable to achieve the same improvements in stability with increasing network
589 layers found by Rasp et al. (2018).

590 **3.3. Overcoming instability in ML emulations and parametrizations**

591 O’Gorman & Dwyer (2018) tackled the instabilities observed in NN-based approaches to subgrid-scale
592 parameterization by employing an alternative ML method; Random Forests (RFs; Breiman 2001; Tibshirani &
593 Friedman 2001). The authors trained a RF to emulate the outputs of a conventional moist convection parameterization

Deleted: -

Deleted: General Circulation Model (

Deleted:)

Deleted: D

Deleted: eural

Deleted: etwork

Deleted: (DNN)

Deleted: D

Deleted: D

Deleted: and DNN-

604 scheme. They then replaced the conventional parameterization scheme with this emulation within a global climate
605 model, and showed that it ran stably and was able to accurately produce climate statistics such as precipitation
606 extremes without needing to be specially trained on extreme scenarios. RFs consist of an ensemble of DTs, with the
607 predictions of the RF being the average of the predictions of the DTs which in turn exist within the domain of the
608 training data. RFs thus have the property that their predictions cannot go outside of the domain for their training data,
609 which in the case of O’Gorman & Dwyer (2018) ensured conservation of energy and nonnegativity of surface
610 precipitation (both critically important features of the moist convection parametrization scheme) were automatically
611 achieved. A disadvantage of this method however is that it requires considerable memory when the climate model is
612 being run to store the tree structures and predicted values which make up the RF.

613 Yuval & O’Gorman (2020) extended on the ideas in O’Gorman & Dwyer (2018), switching from emulation of a single
614 parametrization scheme to emulation of all atmospheric sub grid processes. They trained an RF on a high-resolution
615 three-dimensional model of a quasi-global atmosphere to produce outputs for a course-grained version of the model,
616 and showed that at course resolution the RF can be used to reproduce the climate of the high-resolution simulation,
617 running stably for 1000 days.

618 There are some drawbacks to a RF approach compared to a NN approach however; namely that NNs may provide the
619 possibility for greater accuracy than RFs, and also require substantially less memory when implemented. Given that
620 GCMs are already memory intensive this can be a limiting factor in the practical application of ML parametrization
621 schemes. Furthermore, there is the potential to implement reduced precision NNs on Graphics Processing Units
622 (GPUs) and Central Processing Units (CPUs) which still achieve sufficient accuracy, leading to substantial gains in
623 computational efficiency. Motivated by these considerations, Yuval et al. (2021) trained a NN in a similar manner to
624 how the RF in Yuval & O’Gorman (2020) was trained, using a high resolution aqua-planet model and aiming to coarse
625 grain the model parameters. They overcame the model instabilities observed to occur in previous attempts to use NNs
626 for this process by wherever possible training to predict fluxes and sources and sinks (as opposed to the net tendencies
627 predicted by the RF in Yuval & O’Gorman (2020)), thus incorporating physical constraints into the NN
628 parametrization. The authors also investigated the impact of reduced precision in the NN, and found that it had little
629 impact on the simulated climate.

630 3.4. From aquaplanets to realistic land-ocean simulations

631 All of the studies discussed in this section so far which were tested in a full GCM have used aqua planet simulations.

632 Han et al. (2020) broke away from this trend by developing a Residual NN² (ResNet) based parametrization scheme
633 which emulated the moist physics processes in a realistic land-ocean simulation. Their emulation reproduced the
634 characteristics of the land-ocean simulation well, and was also stable when embedded in single column models.

635 Mooers et al. (2021) represents a subsequent example of an ML emulation of atmospheric fields with realistic
636 geographical boundary conditions, where the authors developed feed-forward NNs to super-parametrize subgrid-scale
637 atmospheric parameters and forced a realistic land surface model with them. Super-parametrization is distinct from
638 traditional parameterization in that it relies on solving (usually simplified) governing equations for subgrid-scale

Deleted: cural

Deleted: etwork

Deleted: D

Deleted: sub grid-scale

Deleted: sub grid-scale

644 processes rather than heuristic approximations of these processes. They employed automated hyperparameter
645 optimization[†] to investigate a range of neural network architectures across ~250 trials, and investigated the statistical
646 characteristics of their emulations. While the authors found that their NNs had a less good fit in the tropical marine
647 boundary layer, attributable to the NN struggling to emulate fast stochastic signals in convection, they also reported
648 good skill for signals on diurnal to synoptic timescales.

649 Brenowitz et al. (2022) sought to address the challenge of emulating fast processes. They used FV3GFS (Zhou et al.,
650 2019; Harris et al., 2021; a compressible atmospheric model used for operational weather forecasts by the US National
651 Weather Service) with a simple cloud microphysics scheme included to generate training data and used this to train a
652 selection of ML models to emulate cloud microphysics processes, including fast phase changes. They emulated
653 different aspects of the microphysics with separate ML models chosen to be suitable to each task. For example, simple
654 parameters were trained with single-layer NNs, while parameters which are more complex spatially were trained with
655 RNNs (e.g., rain falls downwards and not upwards, so it is sequential in timesteps through the atmosphere – a feature
656 which can be represented by an RNN). They then embedded their ML emulation in FV3GFS. They found that their
657 combined ML simulation performed skillfully according to their chosen metrics, but had excessive cloud over the
658 Antarctic Plateau.

659 All of these studies, however, did not test their parameterizations in prognostic long-term simulations.

660 3.5. Testing with prognostic long-term simulations

661 A barrier to achieving stable runs with minimal model drift with ML components is the fact that generic ML models
662 are not designed to conserve quantities which are required to be conserved by the physics of the atmosphere and ocean.
663 Beucler et al. (2019) proposed and tested two methods for imposing such constraints in a NN model; (1) constraining
664 the loss function or (2) constraining the architecture of the network itself. They found that their control NN with no
665 physical constraints imposed performed well, but did so by breaking conservation laws, bringing into question the
666 trustworthiness of such a model in a prognostic setting. Their constrained networks did however generalize better to
667 unforeseen conditions, implying they might perform better under a changing climate than unconstrained models.

668 Chantry et al. (2021b) trained a NN to emulate the non-orographic gravity wave drag parameterization in the ECMWF
669 IFS model (specifically cycle 45R1, ECMWF, 2018) and were able to run stable, accurate simulations out to 1 year
670 with this emulation coupled to the IFS. While the authors note that RFs have been shown to be more stable (e.g.,
671 O’Gorman & Dwyer (2018) and Yuval & O’Gorman (2020), as described above, and Brenowitz et al. (2020)), they
672 chose to focus on NNs since they have lower memory requirements and therefore promise better theoretical
673 performance. The authors assessed the performance of their emulation in a realistic GCM by coupling the NN with
674 the IFS, replacing the existing non-orographic gravity wave drag scheme, and performed 120 hour, 10 day, and 1 year
675 forecasts at ~25 km resolution in a variety of model configurations. The authors showed that their emulation was able
676 to run stably when coupled to the IFS for seasonal timescales, including being able to reproduce the descent of the
677 Quasi-biennial Oscillation (QBO). Interestingly, while the authors initially aimed to ensure momentum conservation
678 in a manner similar to Beucler et al. (2021), they found that this constraint led to model instabilities and that a better

Deleted: D

Deleted: D

681 result was achieved without it. One possible explanation for this is that Beucler et al. (2021) assessed their NNs in an
682 aquaplanet setting. Nonetheless, Chantry et al. (2021b) noted that since their method was not identical to Beucler et
683 al. (2021), improved stability could potentially be achieved by following their method more precisely. The
684 computational cost of the NN emulation developed by Chantry et al. (2021b) was found to be similar that of the
685 existing parametrization scheme when run on CPUs, but was faster by a factor of 10 when run on GPUs due to the
686 reduction in data transmission bottlenecks.

687 The first study to successfully run stable long-term climate simulations with ML parametrizations was Wang et al.
688 (2022a), who extended on the work of Han et al. (2020) by constructing a ResNet to emulate moist physics processes.
689 They used the residual connections from Han et al. (2020) to construct NNs with good nonlinear fitting ability, and
690 filtered out unstable NN parametrizations using a trial-and-error analysis, resulting in the best ResNet set in terms of
691 accuracy and long-term stability. They implemented this scheme in a GCM with realistic geographical boundary
692 conditions and were able to maintain stable simulations for over 10 years in an Atmospheric Model Intercomparison
693 Project (AMIP)-style configuration. This was more akin to a hybrid ML-physics based model than a traditional GCM
694 with ML-based parametrization, because rather than embedding the ResNet in the model code, the authors used a NN-
695 GCM coupling platform through which the NNs and GCMs could interact through data transmission. This is in
696 contrast to the approach employed in the Physical-model Integration with Machine Learning⁴ (PIML) project and
697 Infero⁵, which are both described in Section 3.11. One advantage to this approach noted by the authors is that it allows
698 for a high degree of flexibility in the application of the ML component, however is likely to be less efficient than a
699 fully-embedded ML model, due to the potential for data transmission bottlenecks.

700 3.6. Training with observational data

701 An alternative to using more complex and/or higher resolution models for training data is to train using direct
702 observational data. For example, Ukkonen & Mäkelä (2019) used reanalysis data from ERA5 and lightning
703 observation data to train a variety of different types of ML models to predict thunderstorm occurrence; this was then
704 used as a proxy to trigger deep convection. ML models assessed were logistic regression, RFs, GBDTs, and NNs, with
705 the final two showing a significant increase in skill over convective available potential energy (CAPE; a standard
706 measure of potential convective instability). One of the challenges of accurately reproducing the large-scale effects of
707 convection is correctly identifying when deep convection should occur within a grid cell. The authors proposed that
708 an ML model such as those they assessed could be used as the “trigger function” which activates the deep convection
709 scheme within a GCM.

710 3.7. ML for super parameterization

711 Revisiting the topic of super parameterized subgrid-scale processes introduced above, the use of ML for this approach
712 was investigated in depth by Chattopadhyay et al. (2020). The authors introduced a framework for NN-based super

Deleted: residual DNN[†] (ResDNN)

Deleted: D

Deleted: DNN

Deleted: DNN

Deleted: D

Deleted: sub grid scale

⁴<https://turbo-adventure-f9826cb3.pages.github.io> accessed 7th February 2023

⁵<https://infero.readthedocs.io/en/latest/> accessed 7th February 2023

719 parametrization, and compared the performance of this method against NN-based traditional parametrization (i.e.,
720 based on heuristic approximations of subgrid-scale processes) and direct super parameterization (i.e., explicitly
721 solving for the subgrid-scale processes) in a chaotic Lorenz '96 (Lorenz 1996) system that had three sets of variables,
722 each of a different scale. They found that their NN-based super parameterization outperformed direct super
723 parameterization in terms of computational cost, and was more accurate than NN-based traditional parametrization.
724 The NN-based super parameterization showed comparable accuracy to direct super parameterization in reproducing
725 long-term climate statistics, but was not always comparable for short-term forecasting.

Deleted: sub grid-scale

Deleted: sub grid-scale

726 3.8. Stochastic parametrization schemes

727 A more recent approach to the representation of subgrid-scale processes is via stochastic parameterization schemes,
728 which can represent uncertainty within the scheme. There has been less focus on replacing these schemes with ML
729 alternatives than non-stochastic schemes, however some progress has been made. Krasnopolsky et al. (2013) used an
730 ensemble of NNs to learn a stochastic convection parametrization from data from a high-resolution cloud resolving
731 model. In this case, the stochastic nature of the parametrization was captured by the ensemble of NNs. Gagne et al
732 (2020b) took a different approach, investigating the utility of generative adversarial networks (GANs) for stochastic
733 parameterization schemes in Lorenz '96 (Lorenz 1996) models. In this case, the GAN learned to emulate the noise of
734 the scheme directly, rather than implicitly representing it with an ensemble. They described the effects of different
735 methods to characterize input noise for the GAN, and the performance of the model at both weather and climate
736 timescales. The authors found that the properties of the noise influenced the efficacy of training. Too much noise
737 resulted in impaired model convergence and too little noise resulted in instabilities within the trained networks.

Deleted: sub grid-scale

738 3.9. ML parametrization and emulation for land, ocean, and sea ice models

739 Models of the atmosphere make up one component of the Earth system, however for timescales beyond a few days,
740 simulating other components of the Earth system becomes increasingly important to maintain accuracy. The
741 components which are most often included in coupled Earth system models in addition to the atmosphere are the
742 ocean, sea ice, and the land surface. Reflective of this, ML approaches to parameterization of subgrid-scale processes
743 are not limited to the atmosphere, and progress has been made in the use of ML for land, ocean and sea ice models as
744 well.

Deleted: is

Deleted: .

745 On the ocean modelling front, Krasnopolsky et al. (2002) presented an early application of NN for the approximation
746 of seawater density, the inversion of the seawater equation of state, and a NN approximation of the nonlinear wave-
747 wave interaction. More recently, Bolton & Zanna (2019) investigated the utility of Convolutional Neural Networks
748 (CNNs) for parametrizing unresolved turbulent ocean processes and subsurface flow fields. Zanna & Bolton (2020)
749 then investigated both Relevance Vector Machines[†] (RVMs) and CNNs for parameterizing mesoscale ocean eddies.
750 They demonstrated that because RVMs are interpretable, they can be used to reveal closed-form equations for eddy
751 parameterizations with embedded conservation laws. The authors tested the RVM and CNN parameterizations in an
752 idealized ocean model and found that both improved the statistics of the coarse resolution simulation. While the CNN

758 was found to be more stable than the RVM, the advantage of the RVM was the greater interpretability of its outputs.
759 Finally, Ross et al. (2023) developed a framework for benchmarking ML based parametrization schemes for subgrid-
760 scale ocean processes. They used CNNs, symbolic regression, and genetic programming methods to emulate a variety
761 of subgrid-scale forcings including measures of potential vorticity and velocity, and developed a standard set of
762 metrics to evaluate these emulations. They found that their CNNs were stable and performed well when implemented
763 online, but generalized poorly to new regimes.
764 Focusing instead on sea ice, Chi & Kim (2017) assessed the ability of two NN models: a fully connected NN and an
765 LSTM, to predict Antarctic sea ice concentration up to a year in advance. Their ML models outperformed an
766 autoregressive model comparator, and were in good agreement with observed sea ice extent. Andersson et al. (2021)
767 improved upon this work with their model IceNet, A U-Net ensemble model which produced probabilistic Arctic sea
768 ice concentration predictions to a 6-month lead time. The authors compared IceNet to the SEAS5 dynamical sea ice
769 model (Johnson et al., 2019) and showed an improvement in the accuracy of a binary classification of ice/no ice for
770 all lead months except the first month. Horvat & Roach (2022) used ML to emulate a parameterization of wave-
771 induced sea ice floe fracture they had developed previously, in order to reduce the computational cost of the scheme.
772 When embedded in a climate simulation, their ML scheme resulted in an overall categorical accuracy (accounting for
773 the fact that it was only called where needed) of 96.5%. However, the authors did note that since their ML scheme
774 was trained on present day sea ice conditions, it may have reduced success under different climate scenarios, and they
775 recommend retraining using climate model sea-ice conditions to account for this. Rosier et al. (2023) developed
776 MELTNET, a ML emulation of the ocean induced ice shelf melt rates in the NEMO ocean model (Gurvan et al.,
777 2019). MELTNET consisted of a melt rate segmentation task, followed by a denoising autoencoder network which
778 converted the discrete labelled melt rates to a continuous melt rate. The authors demonstrated that MELTNET
779 generalized well to ice shelf geometries outside the training set, and outperformed two intermediate-complexity melt
780 rate parameterizations, even when parameters in those models were tuned to minimize any misfit for the geometries
781 used. Given the computational cost of sea ice parametrizations is relatively high for the timescales on which sea ice
782 evolution is important (namely, seasonal to climate timescales), and given the promising results in emulating these
783 parametrizations demonstrated in the literature, ML based emulation of these schemes is a strong candidate for
784 inclusion into future dynamical coupled modelling systems.
785 Finally, considering Earth's surface, most of the focus of ML innovations in this context has focused on land use
786 classification (e.g. Carranza-García et al, 2019; Digra et al., 2022) and crop modelling (e.g., Virmodkar et al., 2020;
787 Zhang et al., 2023). The rate of publication of ML applications for land surface models has been slower, however
788 there has nonetheless been steady progress in this space in recent years. Pal & Sharma (2021) presented a review of
789 the use of ML in land surface modelling which provides an excellent primer of the state of the field to that point. They
790 include in their review an overview of land surface modelling components and processes, before reviewing the
791 literature on the use of ML to represent them. They separate their review into attempts to predict and parametrize
792 different variables or aspects of the model, including evapotranspiration (Alemohammad et al., 2017; Zhao et al.,
793 2019; Pan et al., 2020), soil moisture (Pelissier et al., 2020), momentum and heat fluxes (Leufen & Schädler, 2019),
794 and parameter estimation and uncertainty (Chaney et al., 2016; Sawada, 2020; Dagon et al., 2020). They also provide

Deleted: i

Formatted: Font:

Formatted: Font:

796 [a useful summary of the ML architectures that have been used in publications they have discussed. More recently, He](#)
797 [et al. \(2022\) developed a hybrid approach to modelling aspects of the land surface, where a traditional land surface](#)
798 [model was used to optimize selected vegetation characteristics, while a coupled ML model simulated a corresponding](#)
799 [three-layer soil moisture field. The estimated evapotranspiration from this hybrid model was compared to observations](#)
800 [and it was found that it performed well in vegetated areas but underestimated the evapotranspiration in extreme arid](#)
801 [deserts. The ready application of ML to aspects of land surface modelling, and the relative sparsity of publications in](#)
802 [this space suggests that it is a fertile domain for further research and development.](#)

803 **3.10. ML for representing or correcting a sub-component of a parametrization scheme**

Deleted: ¶

804 An alternative method to replacing or emulating an entire parametrization scheme or schemes with ML is to target the
805 most costly or troublesome sub-components of the scheme, and either replace those or make corrections to them.
806 Ukkonen et al. (2020) trained NNs to replace gas optics computations in the RTE-RRTMGP (Radiative Transfer for
807 Energetics and Rapid and accurate Radiative Transfer Model for General circulation models applications-Parallel;
808 Pincus et al., 2019) scheme. The NNs were faster by a factor of 1-6, depending on the software and hardware platforms
809 used. The accuracy of the scheme remained similar to that of the original scheme.
810 Meyer et al. (2022) trained a NN to account for the differences between 1D cloud effects in the European Centre for
811 Medium Range Weather Forecasting (ECMWF) 1D radiation scheme ecRad and 3D cloud effects in the ECMWF
812 SPARTACUS (SPeedy Algorithm for Radiative TrAnsfer through CloUd Sides) solver. The 1D cloud effects solver
813 within ecRad, Tripleclouds, is favored over the 3D SPARTACUS solver because it is five times less computationally
814 expensive. The authors show that their NN can account for differences between the two schemes with typical errors
815 between 20% and 30% of the 3D signal, resulting in an improvement in Tripleclouds' accuracy with an increase in
816 runtime of approximately 1%. By accounting for the differences between SPARTACUS and Tripleclouds rather than
817 emulating all of SPARTACUS, the authors were able to keep Tripleclouds unchanged within ecRad for cloud-free
818 areas of the atmosphere, and utilize the NN 3D correction elsewhere.

819 **3.11. Bridging the gap between popular languages for ML and large numerical models**

820 A common toolset for researchers to develop and experiment with different ML approaches to problems is Python
821 libraries such as pytorch, scikit-learn, tensorflow, keras, etc., or other dynamically-typed, non-compiled languages.
822 In contrast, numerical weather models are almost universally written in statically-typed compiled languages,
823 predominantly Fortran. To make use of ML emulations or parameterizations in the models thus requires that they be:
824 (1) treated as a separate model periodically coupled to the main model (as is done between atmosphere and ocean
825 models for example), or
826 (2) be manually re-implemented in Fortran, or
827 (3) that the pre-existing libraries used are somehow be made accessible within the model code.

828 Wang et al. (2022a; mentioned already above) opted for method 1, developing what could be considered a hybrid ML-
829 physics based model rather than a traditional GCM with ML-based parametrization. In their study, the authors used a

831 NN-GCM coupling platform through which the NNs and GCMs could interact through data transmission. One
832 advantage to this approach noted by the authors is that it allows for a high degree of flexibility in the application of
833 the ML component, however, is likely to be less efficient than a fully-embedded ML model, due to the potential for
834 data transmission bottlenecks. This framework was then formalized by Zhong et al. (2023).
835 There are many examples where method 2 was used, such as Rasp et al. (2018), Brenowitz & Bretherton (2018),
836 Gagne et al. (2019) and Gagne et al. (2020a). The obvious disadvantage of this approach is that every change to the
837 ML model being used requires reimplementing in the Fortran, and if the aim is to test a suite of ML models, this
838 approach becomes untenable. Furthermore, this approach poses greater technical barriers for scientists developing
839 ML-based solutions for numerical model challenges, since they must be sufficiently proficient in Fortran to
840 reimplement models in it, rather than using existing user-friendly Python toolkits.
841 A solution lying somewhere between methods 2 and 3 was developed by Ott et al. (2020), who developed a Fortran-
842 Keras Bridge (FKB) library that facilitated the implementation of Keras-like[†] NN modules in Fortran, providing a
843 more modular means to build NNs in Fortran code. This however did not fully overcome the drawbacks posed by
844 method 2 on its own; implementation of layers in the Fortran is still necessary, and any innovations in the Python
845 modules being used would need to be mirrored in the Fortran library.
846 Finally, method 3 is being tackled by the Met Office in the PIML⁶ project, and by ECMWF with an application called
847 Infero⁷. These projects both seek to develop a framework which can be used by researchers to develop ML solutions
848 to modelling problems in Python, and then integrate them directly into the existing codebase of the physical model
849 (e.g., the Unified model at the UK Met Office). The approach used is to directly expose the compiled code
850 underpinning the Python modules within the physical model code.

851 4. Application of ML for the partial differential equations governing fluid flow

852 The representation and solving of the partial differential equations (PDEs) governing the fluid flow and dynamical
853 processes in the oceans and atmosphere can be considered the backbone of weather and climate models. The solvers
854 used to find solutions to these equations are typically iterative, and must solve the dynamics-governing equations of
855 their model on every timestep and at every grid point. There has been growing interest in using ML to facilitate
856 speedups and computational cost reductions in the preconditioning and execution of these solvers. Preconditioners are
857 used to reduce the number of iterations required for a solver to converge on a solution, and usually do so by inverting
858 parts of the linear problem. Many earlier studies focused on using ML to select the best preconditioner and/or PDE
859 solver from a set of possible choices (e.g. Holloway & Chen, 2007; Kuefler & Chen, 2008; George et al., 2008; Peairs
860 & Chen, 2011; Huang et al., 2016; and Yamada et al., 2018). Ackmann et al. (2020) approached the preconditioner
861 part of the system more directly, using a variety of ML methods to directly predict the pre-condition of a linear solver,
862 rather than using a standard preconditioner. Rizzuti et al. (2019) focused on the solver, using ML to apply corrections

⁶ <https://turbo-adventure-f9826cb3.pages.github.io/> accessed 7th February 2023

⁷ <https://infero.readthedocs.io/en/latest/> accessed 7th February 2023

Field Code Changed

863 to a traditional iterative solver for the Helmholtz equation. Going a step further, a number of studies have used ML to
864 replace the linear solver entirely (Ladický et al., 2015; Yang et al., 2016; Tompson et al., 2017).
865 Representation of the fluid equations in a gridded model poses a challenge because of the inability to resolve fine
866 features in their solution. This leads to the use of course-grained approximations to the actual equations, which aim to
867 accurately represent longer-wavelength dynamics while properly accounting for unresolved smaller-scale features.
868 Bar-Sinai et al. (2019) trained a NN to optimally discretize the PDEs based on actual solutions to the known underlying
869 equations. They showed that their method is highly accurate, allowing them to integrate in time a collection of
870 nonlinear equations in 1 spatial dimension at resolutions $4\times$ to $8\times$ coarser than was possible with standard finite-
871 difference methods.
872 Building on this, Kochkov et al. (2021) developed a ML-based method to accurately calculate the time evolution of
873 solutions to nonlinear PDEs which used grids an order of magnitude coarser than is traditionally required to achieve
874 the same degree of accuracy. They used convolutional NNs to discover discretized versions of the equations (as in
875 Bar-Sinai et al., 2019), and applied this method selectively to the components of traditional solvers most affected by
876 coarse resolution, with each NN being equation specific. They utilized the property that the dynamics of the PDEs
877 were localized, combined with the convolutional layers of their NN enforcing translation invariance[†], to perform their
878 training simulations on small but high-resolution domains, making the training set affordable to produce. An
879 interesting feature of their training approach, which is growing in popularity, was the inclusion of the numerical solver
880 in the training loss function: the loss function was defined as the cumulative pointwise error between the predicted
881 and ground truth values over the training period. In this way, the NN model could see its own outputs as inputs,
882 ensuring an internally-consistent training process. This had the effect of improving the predictive performance of the
883 model over longer timescales, in terms of both accuracy and stability. Finally, the authors demonstrated that their
884 models produced generalizable properties (i.e., although the models were trained on small domains, they produced
885 accurate simulations over larger domains with different forcing and Reynolds number). They showed that this
886 generalization property arose from consistent physical constraints being enforced by their chosen method.
887 An alternative to using ML to discover discretized versions of the PDE equations is to instead use NNs to learn the
888 evolution operator of the underlying unknown PDE, a method often referred to as a DeepONet[†]. The evolution operator
889 maps the solution of a PDE forwards in time and completely characterizes the solution evolution of the underlying
890 unknown PDE. Because it is operating on the PDE, it is scale invariant and so bypasses the restriction of other methods
891 that must be trained for a specific discretization or grid scale. Interest in, and the degree of sophistication of,
892 DeepONets has grown rapidly in recent years (e.g., Lu et al., 2019; Wu & Xiu, 2020; Bhattacharya et al., 2020; Li et
893 al., 2020a; Li et al., 2020b; Li et al., 2020c; Nelsen & Stuart, 2021; Patel et al., 2021; Wang et al., 2021; Lanthaler et
894 al. 2022), to the point where the method is showing promising speedups: $3\times$ faster than traditional solvers in the case
895 of Wang et al. (2021).
896 The application of ML to the solving of PDEs and the preconditioning of PDE solvers has been a fruitful avenue of
897 research to date. It has led to innovations which have proven useful even outside of the immediate field (e.g., Pathak
898 et al. 2022 adapted innovations from DeepONets to use in fully ML-based weather models - this is discussed further
899 in the next Section). This is likely in part because there are many areas of engineering and science which are active in

Formatted: Superscript

900 progressing relevant research, leading to a greater overall pace of innovation. ML-based PDE solvers and
901 preconditioners have not yet been tested in a physical weather and climate model. There are few theoretical reasons
902 this could not occur and, if effective, result in significant computational efficiencies for traditional physical model
903 architectures. This poses an interesting avenue for further research.

904 5. Numerical model replacement/emulation

905 The shift from using ML to emulate or replace parametrization schemes to using ML to replace the entire GCM has
906 been made plausible by the increasing volume of training data available. The focus in this section will be on the
907 challenge of completely replacing a GCM with a ML model.

908 There has been a flurry of activity in the use of ML for nowcasting (e.g. Ravuri et al., 2021), however, since the focus
909 of this review is on weather and climate applications, these studies will not be elaborated on.

910 5.1. Early work – 1D deterministic models

911 Work on the use of ML to predict chaotic time-domain systems initially focused on 1-D problems, including 1-D
912 Lorenz systems (e.g. Karunasinghe & Liong, 2006; Vlachas et al., 2018). Of particular interest is Vlachas et al. (2018),
913 who used Long Short-Term Memory Networks (LSTMs[†]), which are well-suited to complex time domain problems.

914 Convolutional LSTMs (ConvLSTMs), which combine convolutional layers with an LSTM mechanism, were
915 introduced in the meteorological domain by Shi et al. (2015) for precipitation nowcasting. They have since seen wide
916 adoption in other areas (e.g., Yuan et al., 2018; Moishin et al., 2021; Kelotra & Pandey, 2020). Their success in other
917 domains suggests that revisiting their utility for weather and climate modelling could be worthwhile.

Deleted: The recent popularization of convolutional LSTMs, which can also incorporate spatial information, suggests that revisiting the application of LSTMs for the prediction of spatially resolved chaotic systems could prove fruitful.

918 5.2. Moving to spatially extended deterministic ML-based models

919 Replacing a GCM entirely with an ML alternative was first suggested and tested in a spatially-resolved global
920 configuration by Dueben and Bauer (2018), although for this study they only sought to predict a single variable
921 (geopotential height at 500 hPa) on a 6 degree grid. Scher (2018) trained a CNN to predict the next model state of a
922 GCM based on the complete state of the model at the previous step (i.e., an emulator of the GCM). Since this work
923 was intended to be a proof-of-concept, the authors used a highly simplified GCM with no seasonal or diurnal cycle,
924 no ocean, no orography, a resolution of ~625 km in the horizontal, and 10 vertical levels. Nonetheless, their ML model
925 showed impressive capabilities; it was able to predict the complete model state several timesteps ahead, and when run
926 in an iterative way (i.e., by feeding the model outputs back as new inputs) was able to produce a stable climate run
927 with the same climate statistics as the GCM, with no long-term drift (even though no conservation properties were
928 explicitly built into the CNN). Scher & Messori (2019) then extended on this, but continued the proof-of-concept
929 approach. They investigated the ability of NNs to make skillful forecasts iteratively a day at a time to a lead time of a
930 few days for GCMs of varying complexity, and explored a combination of other factors, including number of training
931 years, the effects of model retuning, and the impact of a seasonal cycle on NN model accuracy and stability.

936 Weyn et al. (2019) aimed to predict a limited number of variables, focusing on the NWP to medium range time domain.
937 They trained a CNN to predict 500 hPa geopotential height and 300 to 700 hPa geopotential thickness over the
938 Northern Hemisphere to up to 14-days lead time, showing better skill out to 3 days than persistence, climatology, and
939 a dynamics-based barotropic vorticity model, but not better than an operational full-physics weather prediction model.
940 Weyn et al. (2020) then improved on this significantly, with a Deep U-Net style CNN trained to predict four variables
941 (geopotential height at 500 and 1000 hPa, 300 to 700 hPa geopotential thickness, and 2 m temperature) globally to 14
942 days lead time. A major innovation in this study was their use of a cubed-sphere grid, which minimized distortions
943 for planar convolution algorithms while also providing closed boundary conditions for the edges of the cube faces.
944 Additionally, they extended their previous work to include sequence prediction techniques, making skillful predictions
945 possible to longer lead times. Their improved model outperformed persistence and a coarse resolution comparator (a
946 T42 spectral resolution version of the ECMWF IFS model, with 62 vertical levels and ~2.8 degree horizontal
947 resolution) to the full 14 days lead time, but was not as skillful as a higher resolution comparator (a T63 spectral
948 resolution version of the IFS model with 137 vertical levels and ~1.9 degree horizontal resolution) or the operational
949 subseasonal-to-seasonal (S2S) version of the ECMWF IFS.

950 Clare et al. (2021) tackled a short falling of many of the ML weather and climate models developed to this point,
951 namely that most were deterministic, limiting their potential utility. To address this, they trained a NN to predict full
952 probability density functions of geopotential height at 500 hPa and temperature at 850 hPa at 3 and 5 days lead time,
953 producing a probabilistic forecast which was comparable in accuracy to Weyn et al. (2020).

954 Choosing to focus on improved skill rather than the question of probabilistic vs deterministic models, Rasp & Thuerey
955 (2021) developed a ResNet model trained to predict geopotential height, temperature and precipitation to 5 days lead
956 time and assessed it against the same set of physical models as Weyn et al. (2020). Their model was close to as skillful
957 as the T63 spectral resolution version of the IFS model, and had better skill to the 5 day lead time than Weyn et al.
958 (2020).

959 Keisler (2022) took an ambitious step forward, training a Graph Neural Network[†] (GNN) model to predict 6 physical
960 variables on 13 atmospheric levels on a 1-degree horizontal grid, which the authors claim is ~50-2000 times larger
961 than the number of physical quantities predicted by the models in Rasp & Thuerey (2021) and Weyn et al. (2020).
962 Their model worked by iteratively predicting the state of the 6 variables 6 hours into the future (i.e., the output of each
963 model timestep was the input into the next timestep), to a total lead time of 6 days. The authors showed that their
964 model outperformed both Rasp & Thuerey (2021) and Weyn et al. (2020) in the variables common to all three studies.
965 They suggested that the gain in skill seen over previous studies was due to the use of more channels[‡] of information,
966 and the higher spatial and temporal resolution of their model. Finally, they showed that their model was more skillful
967 than NOAA's GFS physical model to 6 days lead time, but not as skillful as ECMWF's IFS.

968 Lam et al. (2022) also used GNNs to build their ML-based weather and climate model, GraphCast. This model was
969 the most skillful ML-based weather and climate model at the time of writing this review. While the first ML-based
970 weather and climate model to claim to exceed the skill of a numerical model was Pangu-Weather (Bi et al., 2022;
971 described in greater detail in the following subsection), GraphCast exceeded the skill of both the ECMWF
972 deterministic operational forecasting system, HRES, and also Pangu-Weather. Furthermore, Lam et al. (2022) paid

Deleted: Sonderby et al. (2020) took a more targeted approach, developing a NN to produce probabilistic precipitation forecasts to a lead time of 8 hours on a 2 x 2 km resolution grid covering 7000 x 2500 km over the continental United States, with temporal resolution of 2 min and latency (execution time) in the order of seconds. The desired lead time is an input parameter and time-stepping is not used. The focus here was producing rapid high-resolution short-term forecasts of a single key variable.

Deleted: also

Deleted: variables, but focusing

Deleted: ed more

Deleted: Inching slightly closer to being competitive with physical models...

Deleted: DNN

988 particular attention to evaluating their model and HRES against appropriate measures, and included existing model
989 assessment scorecards from ECMWF to evaluate them. GraphCast capitalized on the ability of GNNs to model
990 arbitrary sparse interactions by adopting a high-resolution multi-scale mesh representation of the input and output
991 parameters. It was trained on the ECMWF ERA5 reanalysis archive to produce predictions of five surface variables
992 and six atmospheric variables, each at 37 vertical pressure levels, on a 0.25° grid. It made predictions on a 6-hourly
993 timestep and was run autoregressively to produce predictions to a 10-day lead time. The authors demonstrated that
994 GraphCast was more accurate than HRES on 90.0% of the 2760 variable and lead time combinations they evaluated.

995 5.3. Ensemble generation with ML-based models

996 A common criticism of ML approaches to weather and climate prediction is the difficulty of representing uncertainty,
997 and/or the tails of the distribution of predicted parameters. One common method to represent the range of possible
998 outcomes (including extremes) under different sources of uncertainty is through a well-calibrated ensemble of
999 predictions. There are a growing number of examples where ensemble generation is considered, many of which fall
1000 into the category of full-model replacement.

1001 ▼
1002 Weyn et al. (2021) explored probabilistic ML predictions using an ensemble of NNs similar to the single-member NN
1003 described in Weyn et al. (2020). The authors expanded the number of variables predicted from 4 to 6, and produced
1004 forecasts to 6 weeks lead time - considerably longer than any comparable work at the time of writing this review. They
1005 considered a variety of initial condition perturbation strategies, and explored the impact of model error by varying the
1006 initial values of the model weights during training to create a multi-model ensemble. They used a combination of the
1007 multi-model ensemble generation approach and initial condition perturbations to generate a 'grand ensemble' of 320
1008 members. They used established metrics for ensemble performance such as RMSE-spread plots, and found that the
1009 320-member grand ensemble combining the multi-model ensemble with initial condition perturbations performed only
1010 slightly better than the multi-model ensemble alone at 14 day lead times. The skill of the ensemble mean of the system,
1011 a control member, and the full ensemble were assessed against the same metrics from the ECMWF sub-seasonal to
1012 seasonal (S2S) prediction system. Their grand ensemble had lower skill than the S2S system at shorter lead times, but
1013 was comparable in skill at longer lead times. Their skill assessment used standard probabilistic skill measures such as
1014 continuous ranked probability score and the ranked probability skill score, which are not present in the other studies
1015 discussed in this Section. The next major ML model to be tested in an ensemble mode was FourCastNet, presented by
1016 Pathak et al. (2022), who leveraged the work on DeepONets described in Section 4. In particular, the authors used a
1017 type of DeepONet called a Fourier Neural Operator (FNO). FourCastNet produced predictions of 20 variables
1018 (including challenging-to-predict variables such as surface winds and precipitation) on five vertical levels with 0.25
1019 degree horizontal resolution, and had competitive skill against the ECMWF IFS to 1 week lead time. The high
1020 horizontal resolution of their model enabled it to resolve extreme events such as tropical cyclones and atmospheric
1021 rivers, and the speed of the model facilitated the generation of large ensembles (up to 1,000's of members). ▼

Deleted: Probabilistic

Deleted: and extremes

Deleted: For example, Watson (2022) argued that while there is now an abundance of examples of ML being used for model parameterization schemes, full model replacement, downscaling, and PDE solvers (much of which is covered in this review), there are relatively few examples which address the question of how well ML approaches can reproduce extreme events and statistics.

Deleted: however

Deleted: this is now being

Deleted: Clare et al. (2021) tackled this challenge by training a NN to predict full probability density functions of geopotential height at 500 hPa and temperature at 850 hPa at 3 and 5 days lead time, producing a probabilistic forecast which was comparable in accuracy to Weyn et al. (2020).

Deleted: have also

Deleted: D

Deleted: one

Deleted: and used initial condition perturbation methods and variations in atmospheric representation similar to those used in traditional ensemble prediction to generate the ensemble of DNN predictions.

Deleted: They generated 320-member ensembles (much larger than could be affordably achieved with a physics-based model) and produced forecasts to 6 weeks lead time - considerably longer than any comparable work to date.

Deleted: ML

Deleted: model

Deleted: ¶

Deleted: developed a weather model called FourCastNet,

Deleted: ing

Deleted: This suggests that prediction of very extreme events may be possible. The authors make the ambitious claim that with additional resources and further development, they anticipate that FourCastNet could match the capabilities of current NWP models on all timescales and at all vertical levels of the atmosphere.

1060 The authors explored the potential of their ensemble forecasts by generating a 100-member ensemble from initial
1061 conditions perturbed with Gaussian random noise. They showed that the FourCastNet ensemble mean had lower
1062 RMSE and a higher anomaly correlation coefficient than a single-value prediction at longer lead times (beyond ~3-4
1063 days), although the ensemble mean performed slightly worse than the single value forecast at shorter lead times. The
1064 authors attributed this relative decrease in performance at shorter lead times to the ensemble mean smoothing out fine-
1065 scale features. Unfortunately, the authors did not examine the spread of the ensemble with lead time or evaluate the
1066 model using probabilistic skill metrics (in contrast to Weyn et al., 2021), and while they did consider the capacity of
1067 FourCastNet to predict extremes, they did not do so in an ensemble context.
1068 Hu et al. (2023) improved on the relatively simple ensemble perturbation approach employed by Pathak et al. (2022)
1069 in their model, a Swin (sliding window) Transformer-based Variational Recurrent Neural Network (SwinVRNN).
1070 This model combined a Swin Transformer Recurrent Neural Network (SwinRNN) predictor with a Variational Auto-
1071 Encoder perturbation module. The perturbation module learned the multivariate Gaussian distributions of a time-
1072 variant stochastic latent variable from the training data. The SwinRNN predictor was deterministic, but could be used
1073 to generate ensemble predictions by perturbing model features using noise sampled from the distribution learned by
1074 the perturbation module. Unlike the approach used by Pathak et al. (2022), this strategy ensured that the perturbations
1075 applied at each spatial location in ensemble generation were appropriate for the location and variable in question.
1076 Furthermore, the training strategy employed by Hu et al. (2023) accounted for both the error in the deterministic
1077 predictions and the error in the learned perturbation distribution, effectively optimizing forecast accuracy and
1078 ensemble spread at the same time. The authors assessed both the ensemble spread, and ensemble mean accuracy of
1079 their model, and found that it had a better ensemble spread than simpler alternative ensemble generation strategies.
1080 They also found that it had lower latitude-weighted RMSE than the ECMWF IFS to 5 days lead time for 2m
1081 temperatures and total precipitation. ECMWF data beyond 5 days was not shown, but the SwinVRNN models had
1082 latitude-weighted RMSE values lower than a weekly climatology baseline for three of the four variables shown to 14
1083 days lead time. Bi et al. (2022) achieved a significant milestone with their model Pangu-Weather, the first ML-based
1084 model to perform better than the ECMWF IFS to a lead time of 7 days based on RMSE and Anomaly Correlation
1085 Coefficient (ACC) across several variables including geopotential height and temperature at 500 hPa. While they did
1086 explore the utility of Pangu-Weather for ensemble generation, their approach was more simplistic than that
1087 demonstrated by Hu et al. (2023). Pangu-Weather featured two major innovations over previous contributions to this
1088 space:
1089

1. It used 3D (latitude, longitude and height) input grids trained against 3D output grids. This enabled different
1090 levels of the atmosphere to share information, which was not possible in FourCastNet, in spite of predicting
1091 variables on multiple atmospheric levels, because the levels were treated independently. In contrast, Pangu-
1092 weather adopted a 3D convolutional method that the authors name the 3D Earth-specific transformer
1093 (3DEST), which enabled the flow of information both horizontally and vertically.
- 1094 2. It was made up of a series of models trained with different prediction time gaps. The motivation for this was
1095 that, as noted by the authors, when the goal is to produce forecasts to 5 days (for example), but the timestep
1096 of the basic forecast model is relatively short (e.g. 6 hours), many iterative executions of the model are

Deleted: ¶

Deleted: FourCastNet

Deleted: ,

1100 required, with the errors of each iteration feeding onto the next. A shorter model timestep results in greater
1101 overall errors (due to more iterations being required to reach the final forecast lead time), and a longer model
1102 timestep reduces this error. Motivated by this, the authors trained several versions of their model to predict
1103 to different timesteps on a single iteration. The overall forecast to a given lead time was then constructed
1104 using the longest possible timesteps. For example, for a 7-day forecast, a 24-hour forecast is iterated 7 times,
1105 whereas for a 23-hour forecast, a 6-hour forecast is iterated 3 times, followed by a 3-hour forecast 1 time,
1106 and 1-hour forecast 2 times. The authors noted that this strategy was not effective to multiweek or longer
1107 timescales; they reported that training the model with a 28-day timestep was difficult, for example, and
1108 suggested that more powerful or complex ML methods would be required to achieve this.

1109 As well as the relatively broad measures of RMSE and ACC, the authors assessed the ability of their system to
1110 represent the intensity and track of selected tropical cyclones. They found that Pangu-Weather predicted the tracks of
1111 the cyclones considered with a high degree of accuracy compared to the ECMWF IFS, however it underestimated
1112 cyclone intensity. The authors attributed this to the training data they used (ERA5) also underestimating cyclone
1113 intensity. As noted above, the authors also explored the potential for producing useful ensemble forecasts. To assess
1114 ensemble predictions, they perturbed the initial state of the system with Perlin noise vectors to produce a 100-member
1115 ensemble of forecasts and calculated the RMSE and ACC of the ensemble mean for selected variables. As in Weyn et
1116 al. (2021), the authors noted that the ensemble mean forecasts performed worse than a single deterministic forecast
1117 for shorter lead times (e.g., 1 day), but better for longer lead times. Unfortunately, as with Pathak et al. (2022), Bi et
1118 al. (2022) did not investigate the properties of the spread of the ensemble or assess its skill using standard probabilistic
1119 skill metrics, and their approach to ensemble generation was much simpler than that of Hu et al. (2023).
1120 As already mentioned above, the skill of Pangu-Weather was exceeded by GraphCast, although Lam et al. (2022) only
1121 assessed GraphCast in a deterministic setting. Nonetheless, there is nothing stopping GraphCast from being used to
1122 generate ensemble forecasts in a manner similar to Pangu-Weather. The authors of this review look forward to a more
1123 in-depth intercomparison of the pure ML models in the literature, including an assessment of their performance for
1124 ensemble predictions.
1125 Although the ensemble systems presented in Weyn et al. (2021) and Hu et al. (2023) had lower overall accuracy than
1126 the other models discussed in this section, they still represented the most comprehensive analysis of the behavior and
1127 performance of ensemble ML models (in terms of considering optimal ensemble perturbation strategies, and
1128 quantifying the ensemble behavior) at the time of writing this review. Further investigation into the best methods to
1129 generate and evaluate pure ML model ensembles would be a highly beneficial contribution to the field.

1130 5.4. Moving to more extensible models

1131 As the effectiveness of ML approaches are increasingly demonstrated in the literature, additional factors become clear
1132 in considering these models for both research and application. In a research setting, the ability to readily perform
1133 transfer learning to new problems and reduce training costs will be significant in supporting adoption by other
1134 researchers.

Deleted: intensity., andAs

Deleted:

Deleted: They found that Pangu-Weather predicted the tracks of the cyclones considered with a high degree of accuracy compared to the ECMWF IFS, however it underestimated cyclone intensity. The authors attributed this to the training data they used (ERA5) also underestimating cyclone intensity.

Deleted: T

Deleted: y

Deleted: s

Deleted: The authors did not, however,

Deleted: its utility for probabilistic forecasts or predicting statistical extremes.

Deleted:

Deleted: Nonetheless, there is nothing stopping GraphCast from being used in an ensemble mode, and the assessment of GraphCast presented by Lam et al. (2022) was much more comprehensive and exacting than the assessment of Pangu-Weather presented by Bi et al. (2022).

Deleted: It should be noted that all of the major milestones and high-profile ML models described in this section have relied on reanalysis datasets produced by physics-based models. The provision of higher resolution and higher quality open datasets have the potential to drive progress in this area as much as, if not more than, improvements and further research into ML algorithms.

1162 This need for greater flexibility in both the input data sources and predictive outputs of ML weather and climate
1163 models was recognized by Nguyen et al. (2023), who developed a transformer architecture-based ML model called
1164 ClimaX. This model was designed as a foundational model, trained initially on datasets derived from the CMIP6
1165 (Eyring et al., 2016) dataset, and ~~able to be readily retrained to specific tasks using transfer learning.~~ The authors
1166 demonstrated the skill of ClimaX against simpler ML models, and in some cases a numerical model (ECMWF IFS),
1167 for a variety of tasks including weather prediction, sub-seasonal prediction, climate scenario prediction, and climate
1168 downscaling. The authors showed that ClimaX was able to make skillful predictions in scenarios unseen during the
1169 initial CMIP6 training phase. Furthermore, ClimaX used novel encoding and aggregation blocks in its architecture to
1170 ~~enable greater flexibility in the types of variables used for training, and to reduce training costs when a large number of~~
1171 ~~different input variables were used.~~

Deleted: designed

Deleted: achieve much more affordable training compute costs than other ML weather and climate models such as Graphcast, Pangu-weather and FourCastNet

1172 5.5. Benchmark datasets for ML weather models

1173 Providing open benchmark data for machine learning challenges has been as transformational for the machine learning
1174 field as improved algorithms, the publication of papers, or improvements in hardware.

1175 As the interest and activity in the use of ML as a potential alternative to knowledge-based numerical GCMs has grown,
1176 the need for consistent benchmarks for the intercomparison of ML-based models has become increasingly clear. Rasp
1177 et al. (2020) addressed this need with the introduction of WeatherBench. On this platform, the authors provided data
1178 derived from the ERA5 archive that has been simplified and streamlined for common ML use cases and use by a broad
1179 audience. They also proposed a set of evaluation metrics which facilitate direct comparison between different ML
1180 approaches, and provided baseline scores in these metrics for simple techniques such as linear regression, some deep
1181 learning models and some GCMs. ~~Since the publication of WeatherBench, more benchmark datasets tailored to other~~
1182 ~~domains have been created, including RainBench (de Witt et al., 2020), WeatherBench Probability (Garg et al., 2022),~~
1183 ~~and ClimateBench (Watson-Parris et al., 2022).~~ Weyn et al. (2020) chose datasets and assessment metrics consistent
1184 with WeatherBench to facilitate intercomparison of results. Rasp & Thuerey (2021) directly used the benchmarks
1185 provided by WeatherBench in their assessment. They demonstrated that their model outperformed previous
1186 submissions to WeatherBench, highlighting its value as a tool to allow intercomparability of ML-based weather
1187 models. Other examples of studies using WeatherBench data and analysis methods are Clare et al. (2021) and Weyn
1188 et al. (2021). The parameters of a good benchmark dataset were further elucidated by Dueben et al. (2022), who
1189 provided an overview of the current status of benchmark datasets for ML in weather and climate in use in the research
1190 community and provided a set of guidelines for how researchers could build their own benchmark datasets.

1191 At the time of writing this review, assessments of ML-based models had chiefly (but not exclusively) focused on
1192 simple statistics like globally-averaged RMSE, and not reported in detail on the degree to which they accurately
1193 captured specific processes such as cyclone formation, climate drivers such as the El Nino Southern Oscillation, or
1194 large scale structures such as the jetstreams. A useful contribution from the scientific community would be to better
1195 quantify and articulate a suite of tests and statistics that could form a 'report card' to provide better insight into the
1196 value of new ML models.

1201 [It should also be noted that all of the major milestones and high-profile ML models described in this section so far](#)
1202 [have relied to some degree or another on reanalysis datasets produced by physics-based models. The provision of](#)
1203 [higher resolution and higher quality open datasets have the potential to drive progress in this area as much as, if not](#)
1204 [more than, improvements and further research into ML algorithms.](#)

1205 **5.6. A hybrid approach**

1206 Arcomano et al. (2022) present an approach which straddles the theme of this section and that of the following section
1207 (physics-constrained ML models). Following Wikner et al. (2020), they used a numerical atmospheric GCM and a
1208 computationally-efficient ML method called reservoir computing in a hybrid configuration called Combined Hybrid-
1209 Parallel Prediction (CHyPP). Their hybrid model is more accurate than the GCM alone for most state variables to a
1210 lead time of 7-8 days. They also demonstrate the utility of their hybrid model for climate predictions with a 10-year
1211 long climate simulation, for which they showed that the hybrid model had smaller systematic errors and more realistic
1212 variability than the GCM alone.

1213 **5.7. ML for predicting ocean variables**

1214 More recently, greater attention has been paid to the application of ML to the ocean, particularly for seasonal to multi-
1215 year prediction. Initial work in this space focused on directly predicting key indices such as the NINO 3.4 index. For
1216 example, Ham et al. (2019) trained a CNN to produce skillful El Niño Southern Oscillation (ENSO) forecasts with a
1217 lead time of up to one and a half years. A limiting factor for the application of ML to ocean variables is the lack of
1218 availability of observational data for training. To overcome this, the authors used transfer learning[†] to train their model
1219 first on historical simulations, and then on a reanalysis from 1871 to 1973. Data from 1984 to 2017 was reserved for
1220 validation. Ham et al. (2021) improved on this by including information about the current season in the network inputs
1221 as one-hot vectors[†]. Including this seasonality information led to an overall increase in skill relative to the model in
1222 Ham et al. (2019), in particular for forecasts initiated in boreal spring, a season which is particularly difficult to predict
1223 beyond.

1224 Kim et al. (2022) improved on the performance of the 2D CNNs used in Ham et al. (2019) and Ham et al. (2021) for
1225 predicting ENSO by instead using a convolutional LSTM network with a global receptive field[†]. The move to a larger
1226 (global) receptive field for the convolutional layers enabled the network to learn the large-scale drivers and precursors
1227 of ENSO variability, and the use of a recurrent[†] architecture (in this case LSTM) facilitated the encoding of long-term
1228 sequential features with visual attention[†]. This led to a 5.8% improvement of the correlation coefficient for Nino3.4
1229 index prediction and 13% improvement in corresponding temporal classification with a 12-month lead time compared
1230 to a 2D CNN.

1231 Taylor & Feng (2022) moved from prediction of indices to spatial outputs, training a Unet-LSTM[†] model on ECMWF
1232 ERA5 monthly mean Sea Surface Temperature (SST) and 2-m air temperature data from 1950-2021 to predict global
1233 2D SSTs up to a 24-month lead time. The authors found that their model was skillful in predicting the 2019-2020 El
1234 Niño and the 2016-2017 and 2017-2018 La Niñas, but not for the 2015-2016 extreme El Niño. Since they did not

1235 include any subsurface information in their training data (in contrast to Ham et al. (2019) and Ham et al. (2021), who
1236 included ocean heat content), they concluded that subsurface information may have been relevant for the evolution of
1237 that event.

1238 It is clear from the small number of (but rapidly evolving) studies in this space that there is great promise for the use
1239 of ML for seasonal and multi-year prediction of ocean variables, with many avenues to pursue to achieve potential
1240 skill gains.

1241 5.8. ML for climate prediction

1242 The literature on the use of ML for prediction on seasonal to climate timescales is still relatively sparse compared to
1243 its use for nowcasting and weather prediction. Some examples have been covered in previous sections, such as Weyn
1244 et al. (2021) on subseasonal to seasonal timescales in the atmosphere, and Ham et al. (2019), Ham et al. (2021), Kim
1245 et al. (2022) and Taylor & Feng (2022) on seasonal to multiyear timescales in the ocean. A major cause for this sparsity
1246 is that deep learning typically requires large training datasets, and the available observation period for the earth system
1247 is too short to provide appropriate training data for seasonal to climate timescales in most applications. On the
1248 subseasonal to seasonal end, this may be overcome by including more slowly-varying fields in the training (e.g. ocean
1249 variables), by designing models to learn the underlying dynamics which drive long-term variability, and by including
1250 more physical constraints on the models. On the climate end these same methods could be beneficial, as well as
1251 transfer learning, as is done in Ham et al. (2019), and data augmentation[†] techniques. Additionally, interest is
1252 increasing in the use of ML to predict weather regimes and large-scale circulation patterns, which may prove beneficial
1253 in informing seasonal and climate predictions (Nielsen et al., 2022). [Watson-Parris \(2021\) argued that the differences](#)
1254 [between NWP to multiyear prediction and climate modelling mean that the ML approaches best suited to each can be](#)
1255 [very different. This may also help to explain why the rapid pace of advances in ML based weather models has not](#)
1256 [translated into a similar trend in climate modelling.](#)
1257 Despite this, with the growing maturity of the field of ML for weather and climate prediction, there is every reason to
1258 believe the challenges of prediction on seasonal to climate timescales can be overcome.

Deleted: I

Deleted: W

1259 6. Physics constrained ML models

1260 As has been briefly touched on in previous sections, a promising and increasingly popular method for improving the
1261 performance of ML applications in weather and climate modelling is to include physics-based constraints in the ML
1262 model design (e.g. Karpatne et al., 2017; de Bézenac et al., 2017; Beucler et al., 2019; Yuval et al., 2021; Beucler et
1263 al., 2021; Harder et al., 2022). This can be done through the overall design and formulation of the model, and through
1264 the use of custom loss functions which impose physically-motivated conservations and constraints.

1265 An excellent review of the possible methods for incorporating physics constraints into ML models for weather and
1266 climate modelling, along with 10 case studies of noteworthy applications of these methods, is presented in Kashinath
1267 et al. (2021). The scope of Kashinath et al. (2021) is broad and includes studies not applied directly in the context of

1270 weather and climate modelling, but applicable to it. Rather than repeat the total of this summary here, the reader is
1271 directed to this review.

1272 A class of physics-leveraged ML which has grown rapidly in popularity is Physics Informed Neural Networks
1273 (PINNs). These are discussed in Kashinath et al. (2021), but have also become a very active area of research since the
1274 publication of that review. A more up-to-date review of this class of NNs is presented by Cuomo et al. (2022), along
1275 with a review of other related Physics guided ML architectures.

1276 While PINNs are an exciting and promising new NN architecture, they still face some challenges. For example, they
1277 have had little success simulating dynamical systems whose solution exhibits multi-scale, chaotic or turbulent
1278 behavior. Wang et al. (2022b) attributed this to the inability of PINNs to represent physical causality, and developed
1279 a solution by re-formulating the loss function of a PINN to explicitly account for physical causality during model
1280 training. They demonstrated that this modified PINN was able to successfully simulate chaotic systems such as a
1281 Lorenz system, and the Navier-Stokes equations in the turbulent regime; something which traditional PINNs were
1282 unable to do.

1283 Nonetheless, recent work with PINNs has led to some interesting results for weather and climate simulation: Bihlo &
1284 Popovych (2022) used PINNs to solve the shallow-water equations on a rotating sphere, as a demonstration of their
1285 utility in a meteorological context, and Fuhg et al. (2022) developed a modified PINN to solve interval and fuzzy
1286 partial differential equations, enabling the solving of PDEs including uncertain parameter fields.

1287 7. Other applications of ML and considerations for the use of ML in Weather and Climate Models

1288 Aside from the most active areas of development in the use of ML in weather and climate models discussed in the
1289 sections above, there are a few areas of the literature worth mentioning that are adjacent to the main focus of this
1290 review. These topics are covered in the following subsections.

1291 7.1. Nudging

1292 Rather than replacing a component or components of a GCM with an ML alternative to gain skill improvements, Watt-
1293 Meyer et al. (2021) focused on using corrective nudging to reduce model biases and the errors they can introduce
1294 through feedbacks. The authors used RFs to learn bias-correcting tendencies from a hindcast nudged towards
1295 observations. They then coupled this RF to a prognostic simulation and attempted to correct the model drift with the
1296 learned nudging tendencies. While this simulation ran stably over the year-long test period and showed improvements
1297 in some variables, the errors in others were observed to increase. So far studies in this space seem to be limited to
1298 Watt-Meyer et al. (2021), however this method seems promising, so hopefully interest in developing this approach
1299 further will grow in the future.

1300 7.2. Uncertainty quantification

1301 A common criticism of some ML models such as NNs is that it is difficult to represent the uncertainty of their outputs.
1302 Some examples of studies that have sought to overcome this have already been mentioned in Section 3.8, and there

Deleted: ¶ Object identification within models¶

An alternative to achieving greater model accuracy through increasing resolution of the entire model grid is to develop techniques to identify critical systems and physical phenomena within the model, and embed higher resolution temporary subgrids within the larger GCM to more accurately simulate those processes. A key challenge to overcome to achieve this is automatically identifying key model features. For example, Mudigonda et al. (2017) investigated the feasibility of using a variety of NN architectures to identify storms, tropical cyclones and atmospheric rivers within model data, with promising results. A major limitation of this area of research is the frequent need for labelled datasets of the events being identified, which are currently quite limited. While there are approaches to this problem which utilize unsupervised learning (i.e., learning without an objective function or labelled data), it is harder to achieve a meaningful result this way.

1322 are other examples in the literature (e.g. Grigo & Koutsourelakis, 2019; Atkinson, 2020; Yeo et al., 2021; O'Leary et
1323 al., 2022), however it is nonetheless still a relatively underexplored aspect of ML models for physical systems. Psaros
1324 et al. (2022) suggest that this may be because they are also under-utilized within the broader deep learning community,
1325 and it is thus a developing field that is not universally trusted and understood yet. They also point out that the physical
1326 considerations inherent to ML applied to physical systems often make them more complicated and computationally
1327 expensive than standard ML applications, further disincentivizing the inclusion of uncertainty quantification in an
1328 already complex problem.

1329 Only recently has attention to this aspect of ML become sufficient to motivate the collection of methods into a
1330 consistent framework, a good example of which is the aforementioned Psaros et al. (2022), who presented a
1331 comprehensive review of the methods for quantifying uncertainty in NNs and provided a framework for applying
1332 these methods.

1333 A related topic which is facing similar challenges is the question of explainability of ML approaches; often there is
1334 value in understanding the relative roles and importance of predictors in an ML model, or the relative significance of
1335 different regions of the predictor data. Flora et al. (2022) provide a good overview of approaches to this and compare
1336 their relative drawbacks and benefits.

1337 7.3. Capturing extremes

1338 While there is now an abundance of examples of ML being used for model parameterization schemes, full model
1339 replacement, downscaling, and PDE solvers (much of which is covered in this review), there are relatively few
1340 examples which address the question of how well ML approaches can reproduce extreme events and statistics, both
1341 in terms of the distribution of values predicted in a single-member (i.e., non-ensemble and non-probabilistic) ML
1342 model and in terms of the distribution of predicted outcomes in a probabilistic or ensemble ML model.

1343 Both Pathak et al. (2022) and Bi et al. (2022), introduced in Section 5.2, investigated the ability of their models to
1344 correctly represent extremes, using a similar approach. They divided their test dataset into 50 percentile bins
1345 (distributed logarithmically by Pathak et al. (2022) and linearly by Bi et al. (2022)) between the 90th and 99.99th
1346 percentiles, and computed the relative quantile error between their forecast and ground-truth as a function of lead-
1347 time. Pathak et al. (2022) note that they set their highest percentile bin at 99.99% because of the small sample of
1348 datapoints beyond this percentile making a statistically significant analysis difficult. Both Pathak et al. (2022) and Bi
1349 et al. (2022) found that their models consistently under-forecast extremes to a greater degree than the ECMWF IFS.

1350 Watson (2022) presents a strong argument for the need for a greater focus on the ability of ML weather and climate
1351 models to be able to predict extremes in order for them to meet the needs of users. They present a summary of some
1352 examples of ML models which have sought to predict extreme events according to certain return period definitions.
1353 The example most relevant for this review is Lopez-Gomez et al. (2023), who used a NN with a custom loss function
1354 that preferentially weighted extremes to predict global extreme heat. They found that their custom loss function led to
1355 improved representation of the tails of the distribution (i.e., predictions of extreme heat), and, interestingly, did not
1356 result in any major loss of performance for the middle of the distribution.

1357 The under-prediction of extremes seen in Pathak et al. (2022) and Bi et al. (2022) is consistent with the findings of
1358 Lopez-Gomez et al. (2023), given that neither were not optimized for predicting extremes. These findings all point to
1359 the idea that in order for ML weather and climate models to be able to skillfully predict extreme events, model training
1360 regimes, loss functions and architectures will need to be employed which take into consideration ways to optimize for
1361 these regimes.

1362 **7.4. Object identification within models**

1363 An alternative to achieving greater model accuracy and skill for predicting extremes through increasing resolution of
1364 the entire model grid is to develop techniques to identify critical systems and physical phenomena within the model,
1365 and embed higher resolution temporary subgrids or specialized models within the larger GCM to more accurately
1366 simulate those processes. A challenge to overcome to achieve this is automatically identifying key model features,
1367 since it typically requires a labelled dataset. This requirement can however be avoided, and a variety of both supervised
1368 and unsupervised machine learning approaches to object detection have been demonstrated in the literature.
1369 Mudigonda et al. (2017) were a relatively early example of the application of ML to this challenge. They investigated
1370 the feasibility of using a variety of NN architectures to identify storms, tropical cyclones and atmospheric rivers within
1371 model data, with promising results. Prabhat et al. (2021) provided a valuable resource to the community with their
1372 development of ClimateNet, a labelled open dataset and ML model for the segmentation and identification of tropical
1373 cyclones and atmospheric rivers. This was used by Kapp-Schwoerer et al. (2020) to train a NN to identify and track
1374 these extreme events in Community Atmosphere Model 5 (CAM5; Conley et al. 2012) data. O'Brien et al. (2021)
1375 considered the need for uncertainty quantification in object identification, using a Bayesian approach to build an
1376 atmospheric river detection framework. Finally, Rupe et al. (2023) took a physics-informed approach to object
1377 detection, defining 'local causal states' using speed-of-light causality arguments to identify regions of organized
1378 coherent flow and bypassing the requirement for labelled datasets. They demonstrated the utility of their approach for
1379 the unsupervised identification and tracking of hurricanes and other examples of extreme weather events.
1380 While there are unsupervised learning approaches which have shown value for object detection in weather and climate
1381 data (e.g. Rupe et al., 2023), a major limitation of this area of research is the shortage of labelled datasets for supervised
1382 learning methods, with ClimateNet being an isolated example.

1383 **7.5. GPUs and specialized compute resources**

1384 GPUs and TPUs are specialized hardware which are well suited to highly parallelizable matrix operations, ideal for
1385 solving neural network operations. TPUs have been developed specifically for deep learning applications. Both GPUs
1386 and TPUs are likely to be available on many of the next generation of supercomputers, but much of the current Fortran-
1387 based numerical weather and climate model infrastructure cannot be run on them in their current state. Data
1388 bottlenecks also exist between the GPUs (which have their own on-board memory) and the main memory accessible
1389 to the CPU. While efforts are underway to make numerical and climate models better suited to GPUs, for example
1390 with the development of LFRic (Adams et al. 2019), the new weather and climate modelling system being developed

1391 by the UK Met Office to replace the existing Unified Model (Walters et al. 2017), there is still a long way to go before
1392 entire weather and climate models can be reliably run on GPU or other specialized compute architectures. At the same
1393 time, some neural network designs are aimed squarely at the partial differential equation solving at the core of
1394 numerical methods. Since neural network evaluation utilizes simpler mathematical operations than current PDE
1395 solvers, they offer the prospect of significant computational advantages on non-specialized (i.e., CPU) hardware.

1396 8. Perspectives on machine learning from computer science

1397 This section provides a brief perspective on weather and climate modelling from the computer science domain, and
1398 aims to provide the earth system scientist with a short list of the main relevant innovations in computer science. As
1399 was noted in Section 1, ML models are often regarded as black-boxes, largely because of the design of many prominent
1400 ML systems. In principle, it is not quite right to refer to the trained model as "a machine learning model", in the sense
1401 that the process of training the model is "machine learning", once the model is trained it is definable by a set of
1402 mathematical equations and coefficients, much like any physical, statistical, or theoretical model. Thus the machine
1403 learning refers to the training process, not the model itself. The essence of ML is the level of automation involved.
1404 Even in typical ML models such as large NNs, the model architecture is typically specified manually by the data
1405 scientist or physical scientist involved. The automated derivation of model architecture and composition is not yet
1406 mature for large models, although it is explored through evolutionary programming techniques whereby the learning
1407 of architecture as well as parameterization is automated.

1408 The complex nature of the Earth system means that ML models which seek to emulate it (or subcomponents of it) will
1409 likely also need to be quite complex, and will contain a mixture of ML architectures and algorithms. This is borne out
1410 by the increasing degree of complexity and variety seen in the ML models in the literature reviewed in previous
1411 sections.

1412 A large degree of the current research focus is on very large or deep NNs which rely both on the universal
1413 approximation theorem and practical experimentation to capture a prediction function without needing to explicitly
1414 represent the processes being modeled. In a conceptually similar fashion to how a Fourier decomposition can represent
1415 any wavelike function, the universal approximation theorem establishes that a NN may approximate any function,
1416 subject to its size and the required degree of accuracy (Hornik, Stinchcome and White 1989). Deep learning has been
1417 highly effective in approaching many problems, but many limitations are acknowledged, as evidenced by the current
1418 widespread focus on trustworthy computing and efforts towards explainable ML systems. Some ML models take a
1419 direct approach to modelling the uncertainty of the system being simulated by representing the model state variables
1420 as a probability distribution or degree of confidence. Many contemporary weather and climate model derive their
1421 probabilistic outputs from an ensemble of perturbed members, however an alternative approach is to represent each
1422 part of the belief state[†] of the model as a distribution or likelihood, built up either empirically or by fitting a gaussian
1423 or other known distribution (e.g., Clare et al., 2021).

1424 A timeline of some key innovations in ML is presented in Figure 4. The scale of the timeline is broken between 1956

Deleted: Key papers

Deleted: ,

Deleted: ¶

Deleted: A timeline of some key breakthroughs in ML is presented in Figure XXX Note that the scale of the timeline is broken between 1956 and 1974. Taking that gap in progress into account, it is clear from this visualization that the rate of innovation in ML has increased significantly over the last 80 or so years. This is likely driven by a range of factors including the increasing availability of compute resources suited to ML applications, and the explosion of available data for training. ¶

Deleted: a

Deleted: that either a physical model or a NN could undergo a training cycle (learning) to determine optimal parameter values. The parameters of a neural network are its weights and biases, whereas the parameters of a physical model are physical variables and constants.

Deleted: ¶

Deleted: ¶

As such, the goal of a "ML weather/climate model" (either a full model or an augmented numerical model) will likely be achieved using multiple model types and architectures, in a complex fashion.

Deleted:

Deleted: Current directions in ML are focused on

Deleted: ¶

Deleted: from the computer science domain also

Deleted: have a more fundamental probabilistic or statistical underpinning than typical weather and climate models (... [5])

Deleted:

Deleted: comprise a

Deleted: Whereas a typical

Deleted: or

Deleted: s

Deleted: its

Deleted: ity

Deleted: , an alternative approach could be taken whereby

Deleted: is

Deleted: As such, ML may be applied to statistical moe (... [6])

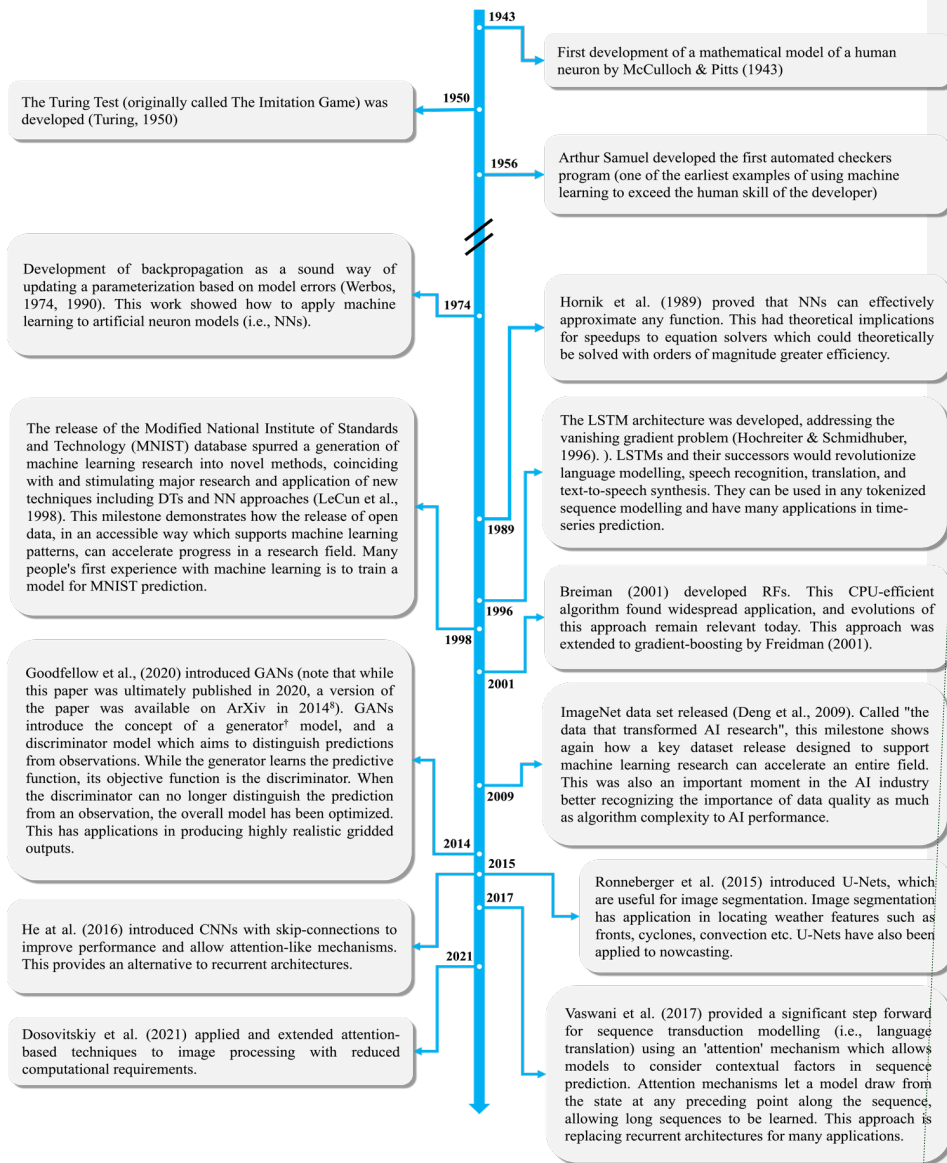
Deleted: ¶

Deleted: ¶

Deleted: breakthroughs

Deleted: XXX

Deleted: (nNote that t



Moved (insertion) [3]

1491

1492

Figure 4: A timeline of key breakthroughs in ML.

1493
1494 and 1974, and Taking that gap in progress into account, it is clear from this visualization that the rate of innovation in
1495 ML has increased significantly over the last 35 or so years. This is likely driven by a range of factors including the
1496 increasing availability of compute resources suited to ML applications, and the explosion of available data for training.
1497 This history shows the degree and rate of research into processing images, text and other sequences based on semantic
1498 understanding of content, but does not demonstrate capturing physical processes as a core element. Advances in the
1499 weather and climate modelling domain have a more explicit goal of properly portraying real physical processes.
1500 Bringing these concepts together promises to uplift capability in both fields.

1501 9. Practical Perspectives on Machine Learning for Weather and Climate Models

1502 A major driver of research into, and improvement of, weather and climate models is increasing the skill of operational
1503 forecast systems, and increasing the accuracy and trustworthiness of climate projections. Therefore, an important
1504 consideration for ML in the context of weather and climate models is the need for it to ultimately be integrated into a
1505 complete predictive system with practical application for forecasting or climate projections.

1506 However, the research findings covered in this review, in spite of being compelling, are yet to make major changes to
1507 operational modelling systems, or standard climate projections.

1508 We have identified three major challenges facing the transition of ML-based innovations into operational settings.
1509 Similar challenges are faced in the context of climate projections, however since these are out of scope for this review
1510 we do not discuss them directly, and instead leave them as a topic for other publications.

1511 The first challenge is the need to assess when a research finding is sufficiently compelling and robust to justify
1512 integration into established operational systems. Since the major function of operational meteorological services is to
1513 inform of future conditions, largely for managing risk or optimizing benefits, a conservative approach is taken to
1514 changing these systems. The utmost premium is put on accuracy, resilience, reliability, and solid scientific foundation,
1515 and many novel research finding require extensive further evaluation and development before they can be considered
1516 ready for inclusion into operational systems. Understanding when to invest this degree of effort in bringing a research
1517 innovation into a major model or scientific configuration upgrade can be difficult.

1518 The second major challenge is establishing the right balance between potentially unwieldy monolithic ML models
1519 which predict all variables of interest, and many smaller limited scope models which each focus on predicting one or
1520 a small number of variables well. The former option is more similar to current dynamical systems, while the latter
1521 option is potentially more easily achievable using an ML approach, but risks becoming difficult to manage due to the
1522 proliferation of small, separate systems. The early effectiveness of limited-purpose ML models provides the ability to
1523 augment existing services without disruption, however, aside from the logistical complexity of many small systems, a
1524 risk associated with this approach is that inconsistencies between predictions may arise from their independent
1525 forecasts, leading to confusion from users and an erosion of trust.

1526 Finally, the third major challenge is how to best monitor and maintain the skill of ML-based systems in a real-time
1527 operational context. Explainability of ML systems is an emerging field, and is not yet sufficiently mature for

Deleted:

Deleted: 4).

Deleted: 80

Deleted: ¶
35¶

Deleted: As such, ML may be applied to statistical models, process-based models, Bayesian models or physical models. Nonetheless, current directions in ML are focused on v... [7]

Deleted: ¶

Deleted: <#>1943: First development of a mathematical model of a human neuron by McCulloch & Pitts (19... [8]

The Turing Test (originally called The Im developed (Turing, 1950)

Development of backpropagation as updating a parameterization based on mo 1974, 1990). This work showed how learning to artificial neuron models (i.e.,)

The release of the Modified National In and Technology (MNIST) database spur machine learning research into novel n with and stimulating major research and techniques including DTs and NN appro: 1998). This milestone demonstrates how data, in an accessible way which support patterns, can accelerate progress in a re people's first experience with machine l model for MNIST prediction.

Goodfellow et al. (2020) introduced GA this paper was ultimately published in the paper was available on ArXiv introduce the concept of a generato discriminator model which aims to disti from observations. While the generator l function, its objective function is the di the discriminator can no longer disting from an observation, the overall model h This has applications in producing high outputs.

He et al. (2016) introduced CNNs with improve performance and allow attentio This provides an alternative to recurrent a

Dosovitskiy et al. (2021) applied and based techniques to image process computational requirements.

Moved up [3]: <#>

Deleted: <#>¶

Deleted: around the world

Deleted: The ...he research findings covered in this review, however ... [9]

Deleted: This section summarizes some practical considerations the research community may wish to be aware of...¶ ... [10]

Deleted: ¶ A ...ince the major function of operational meteorological services is to inform of future conditions, largely for ... [11]

Deleted: ¶ One pathway to adoption of weather and climate models that use ML could be the development of limited-scope m... [12]

1845 application to real-time operational monitoring. Until this changes, the ongoing trustworthiness of operational ML
1846 systems will be difficult to demonstrate. Similarly, online learning in ML weather and climate models is not yet a well
1847 explored research area. The use of online learning is likely to be important for operational ML models to be able to
1848 develop resiliency and maintain good skill over time, so more work will be needed in this area before these models
1849 can see greater uptake in operational systems.

1850 In addition to these major challenges, agencies looking to incorporate ML components into their operational systems
1851 must consider that:

- 1852 • the explainability of ML model errors in the case of poor forecasts that may come under scrutiny,
- 1853 • the robustness of ML models to real-time data issues such as data dropouts or input data degradation must be
1854 established, and
- 1855 • the lack of infrastructure in these agencies to support ML models in an operational setting will need to be
1856 addressed.

1857 Operational development is typically quite incremental, and it is likely that progress will be made in small achievable
1858 steps along the evolving technical frontier. However promising and fascinating as a research direction, full model
1859 replacement with ML alternatives is currently not mature enough for an operational setting. Instead, the authors predict
1860 that the first types of ML systems to be seen in operations will include parameterization scheme replacements and
1861 emulators, solver replacements, super-resolution, new approaches to data assimilation of novel observation sources,
1862 and both pre- and post-processing applications (although of course not all of these have been covered in this review).
1863 It is expected that the research into, and application of, ML methods will represent a growing proportion of weather
1864 and climate model research, with increasingly sophisticated and skillful model components finding their way into
1865 major model releases over the coming years. These components are appealing for both computational and model skill
1866 reasons, and are expected to be highly promising avenues of research.

1867 10. Ethical considerations for Machine Learning for Weather and Climate Models

1868 Not all papers in this review included a discussion of the ethical considerations associated with using machine learning,
1869 nor necessarily touched on what constitutes a sufficiently rigorous verification methodology for machine learning
1870 models. There is a clear relationship between ethical considerations, the explainability of models, and the rigor of
1871 verification applied to ensure that models behave as expected under a variety of conditions (and do not include
1872 unexpected behaviours).

1873 While this review paper does not provide an introduction to AI and ML ethics in general, a brief overview of some
1874 of the important considerations for the application of ML in the context of weather and climate modelling is
1875 provided in this section. Ethical frameworks vary in different cultural and geographical contexts, and for a more
1876 general introduction to the ethical considerations surrounding AI and ML, the reader is directed to the paper
1877 Recommendations on the Ethics of Artificial Intelligence (United Nations Educational, Scientific and Cultural
1878 Organisation (UNESCO), 2022).

- Formatted: Font: 10 pt
- Formatted: Font: Times New Roman, 10 pt
- Formatted: Font: Times New Roman
- Formatted: Font: Times New Roman, 10 pt, Font colour: Text 1
- Formatted: Font: Times New Roman, Font colour: Text 1
- Formatted: Font: Times New Roman, 10 pt, Font colour: Text 1
- Formatted: Font: Times New Roman, 10 pt
- Formatted: Font: Times New Roman, 10 pt, Font colour: Text 1
- Formatted: Font: Times New Roman, Font colour: Text 1
- Formatted: List Paragraph, Bulleted + Level: 1 + Aligned at: 0.63 cm + Indent at: 1.27 cm
- Formatted: Font: 10 pt, Font colour: Text 1
- Deleted: Such an approach could be considered to ent (... [13])
- Deleted: more
- Deleted: , however, I
- Deleted: more
- Formatted: Font colour: Text 1
- Deleted: increments,
- Deleted: system
- Deleted: ¶ (... [14])
- Deleted: applications
- Deleted: ¶
- Deleted: n increasing
- Deleted: ¶ (... [15])
- Deleted: , explainability and reliability concerns
- Deleted: focused on feature commentary on the
- Deleted: involved in
- Deleted: concerns
- Deleted: This
- Deleted: intend to
- Deleted: n
- Deleted: or broad covering of
- Formatted: Line spacing: 1.5 lines
- Deleted:
- Deleted: in general
- Deleted: . The
- Formatted: Font: Italic
- Deleted: published a paper "Recommendations on the (... [16])
- Deleted: ¶ (... [17])

1918 For ML applied to weather and climate modelling, some considerations to ensure sufficient robustness and reliability
 1919 include whether,
 1920 • testing, training and validation data sets are sufficiently representative of the data in general
 1921 • potential causal correlations between testing, training and validation data have been treated correctly
 1922 • trained models have been tested for reliability against adversarial examples
 1923 • data augmentation (e.g. noise addition) has been utilized to enhance model robustness
 1924 • an evaluation of the potential for model drift has been performed
 1925 • the training data is biased in a way which results in ethical unfairness (for example – remote communities
 1926 may not receive equal-skill predictions due to a lack of observational training data in remote areas,
 1927 • the machine learning method is compared to a suitable alternative, such as a known physical model in
 1928 addition to any comparisons to machine learning models or the provision of aggregate statistics
 1929 • the data that has been used has been gathered ethically, and any personal information has been treated
 1930 properly (such as when processing weather reports from individuals)
 1931 • the authors have identified any caveats regarding ethics, reliability, robustness or explainability
 1932 • the authors have investigated the physical realism of the predictions from ML models
 1933 This list is not comprehensive, however. A thorough overview of the explainability, reliability, ethics, and
 1934 verification of ML models in weather and climate has not been covered in prior literature and the field will
 1935 benefit from further work in this area,

1936 **11. Future research directions.**

1937 The already demonstrated and potential future applications for ML in weather and climate modelling are significant
 1938 in number, and identifying the most fruitful avenues for future research can seem overwhelming. A good
 1939 understanding of the current state of the weather and climate modelling field, along with knowledge of the key
 1940 developments in ML research, are required to assess the potential benefits of a given research direction.
 1941 As can be seen from the timeline of machine learning presented in Figure 4, older techniques can prove to be
 1942 relevant many years later, and there are many techniques from computer science which may become relevant for
 1943 contemporary weather and climate modelling problems and research.
 1944 Furthermore, due to the general applicability of many ML approaches, research progresses in one subdomain may
 1945 have implications and benefits for another. For example, DeepONets were developed for, and shown to be
 1946 successful for, solving PDEs, but were adopted by Pathak et al. (2022) for their pure ML model FourCastNet with
 1947 great success.
 1948 To help the reader navigate the myriad research areas where ML for weather and climate modelling could be
 1949 progressed, five categories of future research directions are presented in Figure 5, along with some specific areas of
 1950 research, and benefits that could arise from them.
 1951 These categories are not mutually exclusive – indeed there is overlap between the research areas and benefits
 1952 highlighted in each category (for example, some research focus in Categories 2 and 3 are also applicable to Category

Deleted: A thorough overview of explainability, reliability, ethics and verification has not been covered in prior literature and the field will benefit from further work in this regard. ¶ In the meantime, the authors would like to point to the following particular concerns which are of relevance

Deleted: :

Deleted: Whether

Deleted: Whether

Deleted: Whether

Deleted: Whether

Deleted: Whether

Deleted: Whether

Deleted: Whether

Deleted: Whether

Deleted: Whether

Deleted: Whether

Deleted: Whether

Deleted: Whether

Deleted: Whether

Deleted: ML

Formatted: Indent: Left: 0 cm

Deleted: regard

Deleted: ¶

Deleted: , including additional techniques from machine learning which could be applied

Deleted:

Deleted: ¶

Deleted: often

Formatted: Font: 10 pt

Formatted: Font: 10 pt

Formatted: Line spacing: 1.5 lines

Deleted: stand to

Deleted: have renewed

Deleted: relevance

Formatted: Font: 10 pt

Deleted: earth system science

Formatted: Font: 10 pt

Deleted: ¶

Deleted: ¶

Deleted: ¶

Deleted: ¶

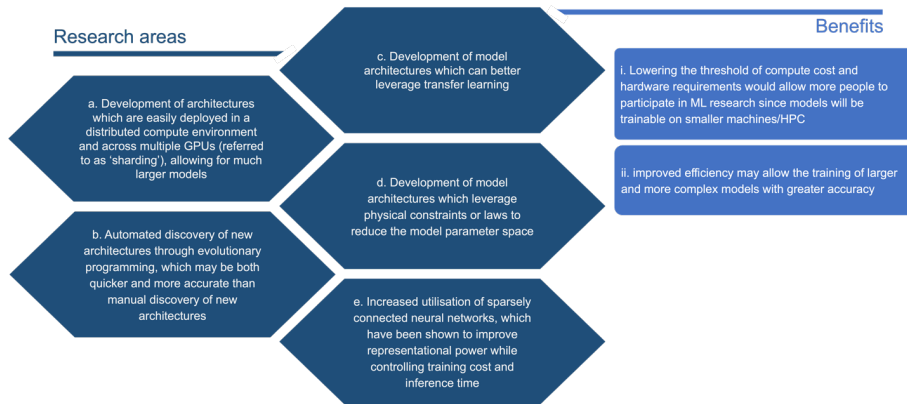
Deleted: ¶

Deleted: usses

1986
1987
1988
1989

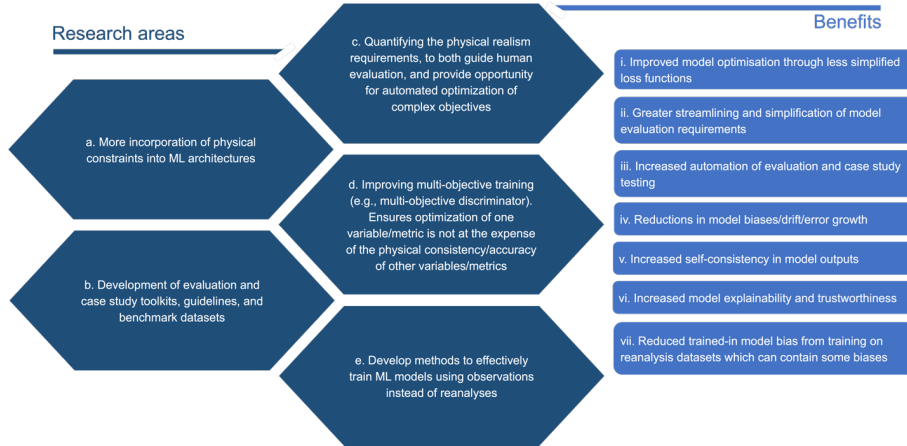
5). The groupings are instead intended to help guide the focus of researchers, and to provide a quick overview of the key topics where the community would most benefit from research progress.

Category 1: Improving training speed and efficiency



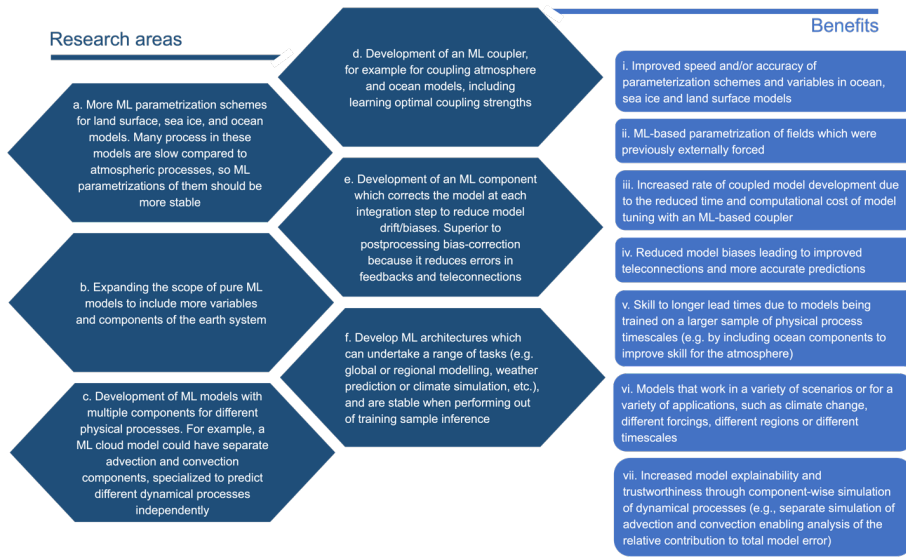
1990

Category 2: Physically consistent/constrained models and evaluation



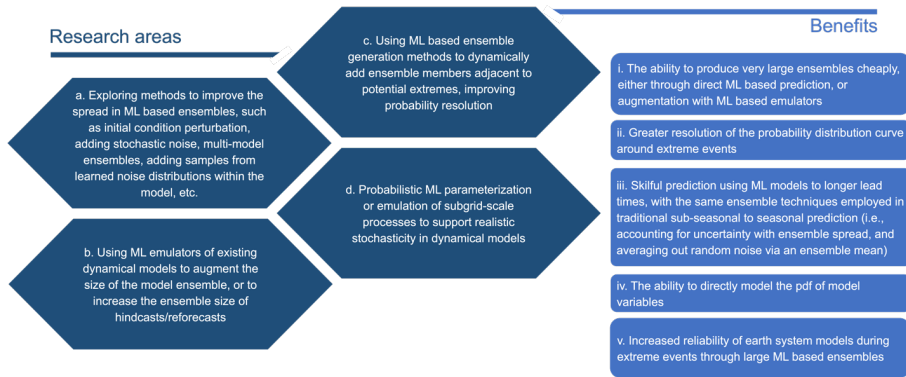
1991

Category 3: Weather and climate modelling domain specific research



1992

Category 4: Probabilistic prediction



1993

Category 5: Trustworthy and explainable systems

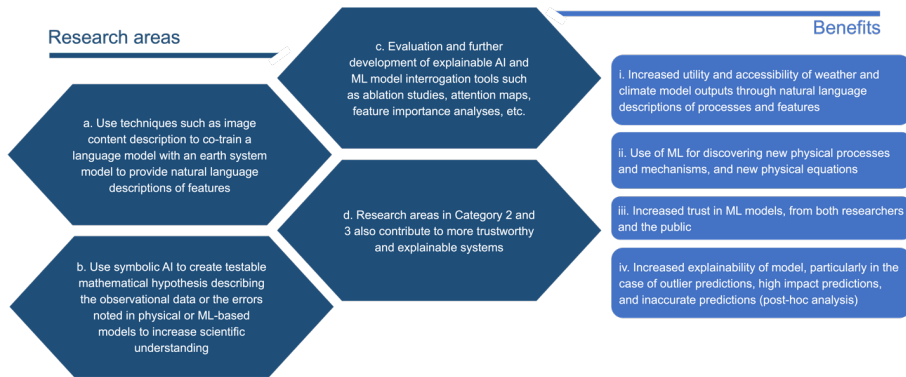


Figure 5: Five categories for future ML research, including suggested research focusses for the community in each category, and potential benefits which could be realized by research and development progress.

Many of the research areas presented are complementary to each other, for example progress in making ML models more affordable to train (Category 1) will increase the utility of ML solutions to a wider community of researchers, and will likely accelerate the rate of progress in the other categories. Progress in the use of physically-informed approaches (e.g. Category 2, area a., or Category 3, area c.) could also lower the training cost of models by reducing the degree of redundancy in the model. On the other hand, approaches such as Category 3, area f., leading to an outcome such as benefit vi. would potentially reduce the demand for more cheaply trainable models, since they could be readily turned to a variety of tasks, saving researchers the need to train their own model from scratch. The research areas and ideas presented here are by no means a comprehensive list. Rather they are intended to be used as a source of inspiration, and the authors of this review are excited to see where the community chooses to focus their efforts in the coming years.

12. Conclusions

In this review we have presented a comprehensive survey of the literature on the use of ML in weather and climate modelling.

We have found that the ML models being most often explored include RFs and NNs, with a high prevalence of FCNNs and CNNs. We have also identified some recent innovations which have proven to be highly effective in the weather and climate modelling space, including DeepONets and variants thereof, Graph NNs, and PINNs.

This review has demonstrated that ML is being successfully applied to many aspects of weather and climate modelling.

We have presented examples from the literature of its application in (1) the emulation and replacement of subgrid-

Formatted: Line spacing: 1.5 lines

Formatted: Font: 10 pt

Deleted: This list provides a list of challenging yet significant future research areas which could provide transformational benefits

Deleted: Table ZZZZ: List of potential research areas

Deleted:

Deleted: ,

Deleted: and DNNs

Deleted: (including ResNets and Deep U-Nets

Deleted:)

Deleted: and

Deleted: sub-grid scale

2033 [scale](#) parametrizations and super-parametrizations, (2) preconditioning and solving of resolved equations, (3) full
2034 model replacement, and (4) a selection of other adjacent areas.

2035 Nonetheless, there are still many [research](#) challenges to overcome, including:

- 2036 • addressing the instabilities excited in physical models due to the inclusion of ML components;
- 2037 • increasing the ease of technical integration (in particular, Fortran compatibility);
- 2038 • memory and computational concerns;
- 2039 • representing a sufficient number of physical parameters and increasing physical and temporal resolution in
2040 ML-based weather and climate model implementations (which currently feature reduced fields and levels
2041 compared to physics-based numerical models);
- 2042 • [moving from a focus on individual parts of the earth system \(i.e., the atmosphere, the ocean, the land surface
2043 etc.\) to tackling the challenges associated with coupled models \(i.e., where models of individual components
2044 of the earth system are coupled together\). Increasingly, operational weather and climate models are coupled
2045 land-atmosphere-ocean-sea-ice models in order to more accurately represent the relevant timescales and
2046 processes in the earth system, and ML modelling efforts need to reflect this;](#)
- 2047 • [more thorough evaluation of the physical realism of ML-based predictions, at various length-scales, across
2048 parameters, and looking at the three-dimensional structures](#)
- 2049 • [Exploring the use of generalized discriminators to augment traditional loss functions in model training \(to
2050 achieve a multivariate generalized objective function\),](#)
- 2051 • the need for more good quality training data; and
- 2052 • the practical challenges of integrating ML components or models into an operational setting.

2053 This list, [together with Section 11](#), provides a set of focus areas for future research efforts. ▲

2054 If the current trend in skill gains in full ML weather and climate models continues, it is possible they will eventually
2055 be considered viable alternatives to traditional numerical models. However, [in the meantime it is likely that ML
2056 components will replace an increasing number of physics-based model components, with models the near-term future
2057 being hybrid ML-physical models. A likely future scenario is one where the best weather and climate models are a
2058 blend of ML and physics-based components, deriving skill from both data driven and physical methodologies.](#)

2059 [Some possible avenues through which increases in ML-based weather and climate model skill might be achieved is
2060 by operating at higher resolutions, resolving more processes which are implicit in the training data, or by undertaking
2061 experiments on synthetic data to address the paucity of real-world data.](#)

2062 Another benefit of ML approaches to weather and climate modeling is the relative computational cheapness of ML
2063 alternatives to current physics-based modelling systems. This has the potential to open the door to experiments that
2064 would not be feasible otherwise. For example, experiments requiring a very large ensemble would be more feasible
2065 with a computationally cheap ML approach.

2066 The literature reviewed here indicates that 'out of the box' ML approaches and architectures are not effective when
2067 used in a weather and climate modelling context. Rather, ML architectures must be adapted to satisfy conservation of

Deleted: multi-objective training strategies which incorporate multiple predicted parameters, physical realism concerns, the capture of statistically extreme values and fine scale structures¶

Formatted: Font: 10 pt, English (US)

Deleted: a more likely scenario

Deleted: is that

Deleted: ¶

2075 energy, represent physically realistic predictions and processes, and maintain good model stability. At the same time,
2076 computational and memory tractability must be maintained.

2077 Advances in the sophistication, complexity and efficiency of ML architectures are being heavily invested in for many
2078 use cases in other disciplines and in the private sector (e.g., condition-action pose estimation, text to video generation,
2079 stable diffusion/text to image, chatbots, facial recognition, semantic image decomposition, etc.). In order to capture
2080 the full benefits of ML for the weather and climate modelling domain, academic and operational agencies will need
2081 to continue to support research in this space. This includes contributing to the research effort through foci such as
2082 those highlighted in Section 11 and in this section, and through addressing the particular challenges facing agencies
2083 interested in the operational and/or realtime deployment of ML based models as the basis for services or the provision
2084 of advice (discussed in Section 9).

2085 Interest and progress in the application of ML to weather and climate modelling has been present for close to 30 years,
2086 and has begun to accelerate rapidly in the last few years. There is good reason to believe that ML as a tool will have
2087 transformational benefits, and offers great potential for further application in weather and climate modelling.

2089 Machine Learning Glossary of Terms

2090 This glossary includes terms which the reader will come across frequently in machine learning literature for the
2091 weather and climate, as well as in machine learning literature generally. Most of these terms are used in this paper
2092 while others support further reading.

2093 **Activation Function.** The function which produces a neuron's outputs given its inputs. Commonly, this includes a
2094 learned bias term which is added to the data inputs before evaluation with a single function to produce the output
2095 value. Examples of the functions used include linear, sigmoid and tanh.

2096 **Adversarial attack.** The deliberate use of malicious data input in a real-world setting intended to cause a
2097 misclassification, underperformance or unexpected behaviours. Examples include emails designed to avoid spam
2098 filters, or images that have been modified to avoid recognition.

2099 **Adversarial example.** A specialised input which results in a misclassification or underperformance of a predictive
2100 model. An example of this concept is an image which has had subtle noise added to it resulting in a copy of that image
2101 which is visually indistinguishable from the original, but which nonetheless causes a misclassification. The term
2102 'adversarial' is used to refer to the way the example fools the model and is not necessarily intended to convey the
2103 sense of malicious intent, although the term is often applied in that fashion. Adversarial examples demonstrate that
2104 machine learning models may be more brittle than expected based on ordinary training data alone. To increase model
2105 robustness, adversarial examples may be generated and added to the training set. Data augmentation techniques such
2106 as flipping, warping and adding noise (any many other techniques) are also used to generate additional training data
2107 to increase robustness and performance.

2108 **Attention mechanism.** A mechanism to allow sequence prediction models to increase the importance of key terms
2109 within that sequence which may be nonlocal and modified in meaning according to the other terms of the sequence.

Deleted: t

Deleted: focusses

Deleted: ¶

Deleted: There are a further set of concerns which are likely to concern those involved in the operational (real-time) deployment of ML based models as the basis for services or the provision of advice, including:¶
The explainability of model errors in the case of poor forecasts that may come under scrutiny¶
Clear guidance on whether models are sufficiently reliable for critical applications¶
How to handle model drift and how to assess model drift in real-time¶
What ongoing training strategies are required for models in a real-time operational setting¶
How ML models may come together as part of an ecosystem of models¶
Robustness to real-time data issues such as data dropouts or input data degradation¶

Deleted: ¶

Deleted: (and potentially highly fascinating, innovative,

Deleted: opportunity

Deleted: **Activation function.** The function which is used to multiply input values, add the bias and produce an output value from an individual node. Examples include linear, sigmoid and tanh.¶

Deleted: complex

2137 **API.** Application Programming Interface. A set of programming functions, methods or protocols by which to build
2138 and integrate applications. APIs may be "web" APIs or imported from software packages in which case they are more
2139 often referred to as libraries.

2140 **Autoencoder.** A neural network architecture which learns to produce a 'code' for an input sequence from which the
2141 original data can be retrieved. The code is shorter than the original input sequence. Applications include data
2142 compression and denoising data.

2143 **Back propagation.** A process of utilising the errors from a prediction to update the weights and biases of a neural
2144 network.

2145 **Batch.** See training batch.

2146 **Batch normalisation.** Data normalisation which aligns the means and variances of input data to a model. For
2147 computational reasons, this is performed separately for each training batch.

2148 **Belief state.** The current state of the world which is believed to be true according to a model. A common architecture
2149 in realtime applications whereby a belief state is updated according to an update function on the basis of new
2150 observations.

2151 **Channel.** An additional dimension to data which is usually not a spatial dimension. Examples include the red, green
2152 and blue intensity images which comprise a colour image. Another example could be to represent both temperature
2153 and wind speed as channels.

2154 **Classification.** A model which attempts to diagnose or predict the category, label, class or type that an example falls
2155 within.

2156 **Climatology.** Refers to the usual past conditions for a location at a time of year. Usually calculated by temporal mean
2157 across years of a dataset, for a given time interval within those years (e.g., for a dataset of monthly mean values
2158 spanning all months of all years from 1990 to 2020, the monthly mean climatology would be obtained by averaging
2159 across all the Januarys from each year, all the Februarys, etc., to obtain an "average January", an "average February",
2160 etc.). Climatologies are often used in the same manner as persistence as a baseline prediction against which to measure
2161 a predictive model. For example, a model predicting a value for January could be compared to the climatological
2162 monthly mean value for January. This helps answer the question "is my model a better source of information than
2163 using the average past conditions from this time of year?".

2164 **Connectome.** The connections between nodes in a neural network. Examples include fully-connected, partially-
2165 connected, skip-layer connections, recurrent connections and others. The 'wiring diagram' for the network.

2166 **Convolutional neural network.** A neural network architecture commonly applied to images which utilises a
2167 convolutional (spatially connected) kernel applied in a sliding window fashion with a narrow receptive field to
2168 encourage the network to generalise from fine scale structure to higher levels of abstraction.

2169 **Data augmentation.** The practice of modifying input data in supervised learning to produce additional examples.
2170 This can make networks more robust to new inputs and address issues of brittleness to adversarial examples. An
2171 example of data augmentation is using rotated or reflected versions of the same image as independent training samples.

2172 **Data driven.** A generalised term used to indicate a primary reliance or dependence on the collection or analysis of
2173 data. Used in contrast to process driven or theory driven.

Deleted: deep

Deleted: small

Deleted: across the data.

Deleted:

2178 **Decision tree.** A tree-like, or flowchart-like, branching model representing a series of decisions and their possible
2179 consequences. Each internal node represents a 'test' (i.e. decision threshold) and each leaf node represents a class label
2180 or collection of possible outcomes.

2181 **Deep NN.** A neural network with many layers. Deeper, thinner networks have [generally been more popular in recent](#)
2182 [times](#) than wider, shallower ones [but this is not always the case \(see e.g. Zagoruyko & Komodakis, 2016\)](#).

2183 **DeepONet.** A neural network architecture relying on universal approximation theorem to train a neural network to
2184 represent a mathematical operation (the operator), such as a partial differential equation or dynamic system.

2185 **Discriminator model.** A model which distinguishes or discriminates between synthetic data and real-world
2186 observations. Often used in conjunction with a generator. In this case, the overall goal is to produce a generator which
2187 is capable of fooling the discriminator, producing highly realistic images. This process is used in Generative
2188 Adversarial Networks.

2189 **Dropout layer.** A neural network layer which is only partially connected, often with a stochastic dropout chance. This
2190 has been shown experimentally to improve neural network robustness in many architectures by reducing overfitting.

2191 **Epoch.** A single complete training pass through all available training data, e.g. learning from all samples, or learning
2192 from all mini-batches, according to the training strategy. Multiple training epochs will typically be utilised although
2193 alternative strategies do exist.

2194 **Feed-forward network.** A neural network composed of distinct 'layers', where the outputs of one layer never feed
2195 back into earlier layers. This avoids the needs for any iterative solver approaches and results in a very computationally
2196 efficient 'forward pass'.

2197 **Generative adversarial network.** A two-part neural network architecture comprising a generator and a discriminator,
2198 which are co-trained to produce realistic outputs which are hard to distinguish from real-world data. The discriminator
2199 replaces the traditional loss function.

2200 **Generator model.** A model which produces a synthetic example of a particular class, such as a synthetic image or
2201 synthetic language. Examples include language or image generation. These are used as part of Generative Adversarial
2202 Networks among other applications.

2203 **Global receptive field.** Where every part of the input region can influence or stimulate a response in a model (e.g. a
2204 fully-connected neural network).

2205 **GPU.** Graphical Processing Unit. A hardware device specialised for fast matrix operations, originally created to
2206 support computer graphics, particularly for games.

2207 **Gradient boosted decision tree.** Also referred to as extreme gradient boosting. A random forest architecture which
2208 combines gradient boosting with decision tree ensembles.

2209 **Gradient boosting.** An approach to model training where each additional ensemble member attempts to predict the
2210 cumulative errors of previously trained members.

2211 **Graph neural network.** A class of neural networks designed to process data which is described by a graph (or
2212 tree/network) data structure. [See Scarselli et al. \(2008\), Kipf & Welling \(2016\), and Battaglia et al. \(2018\) for more](#)
2213 [information and examples](#).

Deleted: neural network

Deleted: proven

Deleted: easier to train

Deleted: <https://arxiv.org/abs/1605.07146>

Deleted: .

Deleted: deep

2220 **Hidden layer.** A layer which is intermediate between the input layer and the output layer of a network or tree structure.
2221 Hidden layers may be used to encode 'hidden variables' which are latent to a problem but not able to be directly
2222 observed.
2223 **Hierarchical temporal aggregation.** A mechanism of composing neural networks which are trained for different lead
2224 times to produce an optimal prediction at all time horizons.
2225 **Hierarchical temporal memory.** Fundamentally different to hierarchical temporal aggregation. A complex deep
2226 learning architecture which uses time-adjacency pooling.
2227 **Hyperparameter.** A parameter which is not derived via training. Examples include the learning rate and the model
2228 topology.
2229 **Hyperparameter search (or Hyperparameter optimization).** The process of determining optimal hyperparameters.
2230 This term may also be used to encompass the model selection problem. This process is automated in some cases.
2231 **Input layer.** A layer which is composed of input nodes. Typically machine learning models will have one input layer
2232 at depth zero (i.e. with no preceding layers) and no input nodes at greater depths.
2233 **Input node.** A node which represents an input or observed value.
2234 **K-fold cross-validation.** A process of changing the validation and test data partitions during different iterations of
2235 training. This allows more of the training and validation data to be used while minimising overfitting. Some definitions
2236 include test data in this process but that is not ideal as the final test is no longer statistically independent.
2237 **Keras.** A streamlined API for creating neural networks, integrated with Tensorflow. Originally built on the Theano
2238 framework for general mathematical evaluation. PyTensor and Aesara are related packages.
2239 **Kernel trick.** For data sets which are not linearly separable, first multiplying the data by a nonlinear function in a
2240 higher dimension can result in a linearly separable higher-dimensional data set to which a simpler method can be used
2241 to model the data.
2242 **Knowledge based systems.** A broad term from artificial intelligence meaning a system which that uses reasoning and
2243 a knowledge base to support decision making. Knowledge is represented explicitly and a reasoning or inference engine
2244 is used to arrive at new knowledge.
2245 **Layer.** In tree or feed-forward network structures (e.g. decision trees and feed-forward neural networks), a layer refers
2246 to the set of nodes at the same depth within a network.
2247 **Leaf node.** Aka output node. A node which does not have any child nodes.
2248 **Long short term memory network.** A recurrent neural network architecture which processes sequences of tokens
2249 utilising a 'memory' component which can store information from tokens early in a sequence for use in prediction of
2250 tokens much later in a sequence. Typical applications include language prediction and time-series prediction of many
2251 kinds.
2252 **Loss function** (also known as target function, training function, objective function, penalty score, error function,
2253 heuristic function, minimisation function). A differentiable function which is well-behaved, such that smaller values
2254 represent better model performance and larger values represent worse performance. An example would be the root-
2255 mean-squared-error of a prediction compared to the truth or target value.

Deleted: , may constitute layers of a deep neural network, or may have other purposes

Deleted: s

2259 **Mini batch.** A subset or 'mini batch' of the training data. Utilised for multiple reasons, including computational
2260 efficiency and to reduce overfitting. Aggregate error over a mini-batch is be learned rather than per-sample errors.
2261 This is the typical contemporary approach. See also training batch for in-depth discussion.

2262 **Neural network.** A composition of 'input nodes', 'connections', 'nodes', 'layers', 'output layers' and 'activation
2263 functions' which are capable of complex modelling tasks. Originally designed to simulate human neural functioning
2264 and subsequently applied to a range of applications.

2265 **Node. Aka vertex.** A small data structure in a network, tree or graph structure which is connected by edges. A node
2266 may represent a real-world value (such as a location) or an abstract value (such as in a neural network), or a decision
2267 threshold (such as in a decision tree).

2268 [Normalisation. A technique applied in many areas of mathematics, science and statistics which is also very important](#)
2269 [to machine learning and neural networks. In a general sense, this refers to expressing values within a standard range.](#)
2270 [Very often, the range of expected values is mapped onto the range 0 to 1, to allow physical variables with different](#)
2271 [measurement units to be compared on equal scale. Such normalisation may be linear or nonlinear, according to a](#)
2272 [simple or more complex function, and either drawn from known physical limits or from the variation observed in the](#)
2273 [data itself.](#)

2274 **One-hot vector.** A vector of 1s and 0s, in which only one bit is set to 1. Typically produced during the first step in
2275 machine learning for language processing to create a word or feature embedding in a process called tokenisation or
2276 encoding. The length of the vector is commonly equal to the number of categories or symbols.

2277 **Output layer.** A layer which comprises the leaf nodes or output nodes of a tree or network.

2278 **Perceptron.** A single-layer neural network architecture for supervised learning of binary classification. Originally
2279 built as an electronic hardware device encoding weights with potentiometers and learning with motors. A multi-layer
2280 perceptron is the same thing as an ordinary neural network.

2281 **Persistence.** Refers to the practice of treating some past observation or reanalysis (usually immediately prior to the
2282 starting point of the prediction period) as the future prediction and "persisting" this one state forward to every
2283 prediction lead time. The predictive model is then compared to this persistence prediction, essentially assessing the
2284 performance of the model against a steady state prediction. This, along with climatology, is often used as a baseline
2285 or bare minimum prediction to beat (i.e., a prediction better than persistence could be considered skilful vs
2286 persistence). This answers the question "is my model a better source of information than using what happened just
2287 before now?".

2288 **Physically-informed machine learning. Also known as physics-informed machine learning.** Machine learning is
2289 considered physically informed when some aspect of physics is included in any way. Examples include adding a
2290 physical component to the loss function (e.g. to enforce conservation of physical properties) or using an activation
2291 function with physically realistic properties.

2292 **Predictive step, forward pass, evaluation.** The process of calculating a model prediction from a set of input
2293 conditions. Distinct from the training phase or back-propagation step.

2294 **PyTorch.** A widely adopted framework for neural networks in Python.

2295 **Random forest.** An architecture based on decision tree ensembles where each decision tree is initialised semi-
2296 randomly and an average of all models is used for prediction. This is typically more accurate than a single decision
2297 tree but less accurate than a gradient-boosted decision tree and so is now less-used. The term random forest is still
2298 commonly used when in fact the implementation is a gradient boosted decision tree.

2299 **Receptive field.** The size or extent of a region in the input which can influence or stimulate a response in a model,
2300 e.g. the size of a convolutional kernel, the size of a sliding window

2301 **Rectified Linear Unit (ReLU).** An activation function commonly used in DNNs. Defined as $\max(0, X)$. This function
2302 is used as it is computationally cheap and avoids problems of vanishing gradients.

2303 **Recurrent network.** A neural network which does pass the output from nodes of the network back into the input of
2304 others. Infinite recurrence is avoided by setting a specific number of iterations for the recurrence. These are often
2305 depicted in diagrams as separate layers but the implementation is through internal recurrent connections.

2306 **Regression.** A model which attempts to diagnose or predict an exact value by statistically relating example input
2307 values to desired values.

2308 **Relevance vector machine.** A sparse Bayesian model utilising the kernel trick in similar fashion to a support vector
2309 machine.

2310 **Representation error.** Error which is introduced due to the inexactness of representing the real world in the model
2311 belief state. Examples may include topography smoothing, point-to-grid translations, model grid distortions near the
2312 poles, or the exclusion of physical characteristics which are not primary to the model.

2313 **Residual neural network (ResNet).** A very influential and innovative convolutional NN architecture which uses a
2314 similar concept to gradient boosting. Each layer of the deep network is taken to predict the residual error from the
2315 previous layers, with skip-connections from earlier layers allowing the training to occur without the issue of vanishing
2316 gradients.

2317 **Sample.** A single training example (e.g. a row of data).

2318 **Scale invariance.** A feature of a system, problem or model which means the results and behaviour are the same at any
2319 scale (e.g., the behaviour does not change if the inputs are multiplied by a common factor).

2320 **Scikit-learn.** A popular Python library for machine learning which extends the SciPy framework.

2321 **Sharding.** Refers to dividing the training of a neural network across multiple GPUs or nodes. This can be done using
2322 data sharding, whereby each GPU or node trains on a subset of the data to allow training parallelism, or model sharding
2323 where a single model is partitioned across multiple GPUs to allow a larger neural network than could be allocated in
2324 memory on a single GPU. One example could be assigning a small number neural network layers to each GPU which
2325 could then work in sequence to operate on a very large network.

2326 **(Stochastic) Gradient descent.** An algorithm by which a neural network is trained using increasingly fine-scale
2327 adjustments to optimise the accuracy of network prediction. Utilised to find the local minimum of a differentiable
2328 function.

2329 **Supervised learning.** Machine learning is considered 'supervised' when the data is labelled according to a category
2330 or target value. Classification data have an explicit labelled category. Regression data have an explicit value which is
2331 being predicted for.

Formatted: Font: Bold

Deleted: deep

Deleted: deep learning

Formatted: Font: Bold

Formatted: Font: Not Bold

2334 **Support vector machine.** A classification model based on finding a hyperplane to separate data utilising the kernel
2335 trick.

2336 **Tensor.** Can be considered as a dense multi-dimensional array or matrix.

2337 **Tensorflow.** A widely adopted framework for neural networks in Python.

2338 **Test/train/validate split.** Available data is split into three portions. The training data is evaluated and used to update
2339 model weights. Validation data is evaluated during training and may be used for hyper-parameter search or to guide
2340 the researcher. Test data is independent (typically well-curated) data used for gold standard evaluation. In reality,
2341 validation data is sometimes used as test data, but this is not good practice. There are many considerations for
2342 test/train/validate splitting, such as statistical independence, representation of all classes, and bias. It is important to
2343 consider what the model is generalising "from" and "to", and ensuring appropriate examples are present in the training
2344 data and appropriate examples are reserved for validation and test.

2345 **Token.** Tokenisation the process of mapping a symbolic or categorical sequence to a numerical representation which
2346 is suited to a sequence-based machine learning model. Commonly, a vector representation will be utilised for the token
2347 form. In language processing, either characters or words may be represented as tokens depending on the approach.

2348 **Top Hat function.** A filter or function which has a rectangular shape resembling the cross-section of a top hat. One
2349 of the simplest functions used for convolutional operations, it can be defined as one constant value in a given bounded
2350 range, and another smaller constant value outside that range.

2351 **TPU.** Tensor Processing Unit. A hardware device specialised for artificial intelligence and machine learning
2352 applications, in particular neural network operations.

2353 **Training batch (or simply batch).** Multiple definitions apply and the use the term has evolved over time. Originally
2354 used in the context of learning from offline or saved historical data as opposed to online or realtime novel data. In this
2355 definition, the training batch is the saved data and refers to the whole training set. For example, a robot exploring a
2356 new environment in real-time must use an online learning technique and could not utilise batch training to map the
2357 unseen terrain. In more recent use, particularly in the areas of neural network learning, the offline saved data may be
2358 split into one or more batches (subsets). If one batch (the batch is the entire training set) is used, the aggregate errors
2359 for the entire training set are used to update the model weights and biases, and the learning algorithm is called batch
2360 gradient descent. If each example is presented individually, this is called online training (even when historical saved
2361 data is being used), the weights and biases are updated for from each individual example, and the algorithm used is
2362 stochastic gradient descent. If the data is divided into multiple batches, this is often referred to equivalently as mini
2363 batches. The weights and biases are aggregated over each mini batch. This is the most common contemporary
2364 approach, as it reduces overfitting and is a good balance of training accuracy, avoiding local minima, and
2365 computational efficiency.

2366 **Transfer learning.** The process of training a model first on a related problem, and then conducting further training
2367 on a more specific problem. Examples could be training a model first in one geographical region and then in another;
2368 or training first at a low resolution then subsequently at a high resolution. This is frequently done to reduce training
2369 computation cost for similar problems by re-using the trained weights from a well-performing source model, or to
2370 overcome a problem of limited data availability by using multiple data sources.

Formatted: Font: Bold

2371 **Transformer network.** A token-sequence architecture which is capable of handling long-range dependencies.
2372 Initially applied to language processing, it has found effective application in image processing as an alternative to
2373 convolutional architectures.

2374 **Translation invariance.** A feature of a system, problem or model which means the results and behaviour are the same
2375 after any spatial translation (i.e., the behaviour does not change if the inputs are shifted spatially to a new location).

2376 **U-Net.** A type of convolutional neural network developed for biomedical image segmentation which has found broad
2377 application. In the contracting part of the network spatial information is reduced while feature information is increased.
2378 In the expanding part of the network, feature information is used to inform high-resolution segmentation. The name
2379 derives from the diagrammatic shape of the network forming a "U".

2380 **Unsupervised learning.** Machine learning is considered 'unsupervised' when data is unlabelled. Examples include
2381 clustering, association and dimensionality reduction.

2382 **Vanishing Gradient.** At the extremes, nonlinear functions used to calculate gradients can result in gradient values
2383 which are effectively zero. These small or zero values, once present in the weights and biases of a neural network, can
2384 entirely suppress information which would in fact be useful, and result in a local minima from which training cannot
2385 recover. This is particularly relevant to long token-series when long-distance connections are relevant. A variety of
2386 techniques including alternative activation functions, training weight decay, skip connections and attention
2387 mechanisms may each or all be utilised to ameliorate this issue.

2388 **Weights and biases.** The parameter values for each neuron which represent the weighting factors to apply to the input
2389 values, plus an overall bias value for the node.

2390 **XGBoost.** A popular Python library for gradient boosted decision trees.

2391 **Appendix A: Table Summary of Model Architectures cited in this paper.**

2392 This table includes all references from this review except for: seminal ML papers that are on new ML methods (e.g., foundational
2393 ML papers), review papers, any paper cited that concerns a topic which is out of scope (e.g., nowcasting), and any other paper
2394 which does not present a new method directly applicable to weather and climate modelling.

Formatted: Font: 9 pt

Author(s)	Year	Category	Approach
Ackmann et al	2020	Fully connected NN	Preconditioner
Alemohammad et al	2017	Fully connected NN	Variable estimation
Andersson et al	2021	Convolutional NN	Prediction
Arcomano et al	2022	Reservoir computing	Alongside-model bias corrector
Atkinson	2020	Baysean type NN	PDE solver
Bar-Sinai	2019	Convolutional NN	PDE solver
Battaglia et al	2018	Graph NN	Method paper
Beucler et al	2019	Physics Informed NN	Convective paramterisation
Beucler et al	2021	Physics Informed NN	Convective paramterisation
Bhattacharya et al	2021	Fully connected NN	PDE solver
Bi et al	2022	Mixed/Custom NN	Pure ML atmospheric model

Bihlo & Popovych	2022	Physics Informed NN	PDE solver
Bolton and Zanna	2019	Convolutional NN	Parametrization
Brenowitz & Bretherton	2018	Fully connected NN	Parametrization
Brenowitz & Bretherton	2019	Fully connected NN	Parametrization
Brenowitz et al.	2020	Fully connected NN	Parametrization
Brenowitz et al	2020	Decision tree-based, Fully connected NN	ML model intercomaprison
Brenowitz et al	2022	Recurrent NN	Parametrization
Chaney et al	2016	Decision tree-based	Interpolation
Chantry et al	2021	Fully connected NN	Parametrization
Chattopadhyay et al	2020	Fully connected NN, Recurrent NN	Super parametrization
Chevallier et al	1998	Fully connected NN	Parametrization
Chi & Kim	2017	Fully connected NN, Recurrent NN	Prediction
Clare et al	2021	ResNet	Emulation (probabilistic)
Dagon et al	2020	Fully connected NN	Emulation
de Bézenac et al	2017	GAN	Prediction, model evaluation
Deuben and Bauer	2018	Fully connected NN	Replacement
Flora et al	2022	Decision tree-based, Logistic regression	Assessment of explainability techniques
Fuhg et al	2022	Physics Informed NN	PDE solver
Gagne et al	2019	Decision tree-based	Parametrization
Gagne et al	2020	GAN	Parametrization (probabilistic)
Gagne et al	2020	GAN, Fully connected NN	Parametrization
George et al	2008	Mixed/Custom non-NN	Preconditioner
Gettelman et al	2021	Fully connected NN	Emulation
Ham et al	2019	Convolutional NN	Prediction
Ham et al	2021	Convolutional NN	Prediction
Han et al	2020	ResNet	Parametrization
Harder et al	2022	Fully connected NN	Emulation
He et al	2022	Decision tree-based	Parametrization
Holloway & Chen	2007	Fully connected NN	Preconditioner and PDE solver selection
Horvat & Roach	2022	Fully connected NN	Parametrization
Hu et al	2023	Mixed/Custom NN	Pure ML atmospheric model
Huang et al	2016	SVM	Preconditioner
Kapp-Schwoerer et al	2020	Convolutional NN	Semantic segmentation
Karunasinghe & Liong	2006	Fully connected NN	Chaotic timeseries prediction
Keisler	2022	Graph NN	Replacement
Kim et al	2022	Mixed/Custom NN	Prediction
Kochkov et al	2021	Convolutional NN	PDE solver

Krasnopolsky et al	2002	Fully connected NN	Emulation
Krasnopolsky et al	2005	Fully connected NN	Emulation
Krasnopolsky	2013	Fully connected NN	Parametrization (probabilistic)
Kuefler & Chen	2008	Mixed/Custom non-NN	Linear system solver
Ladický et al	2015	Decision tree-based	PDE solver
Lam et al	2022	Mixed/Custom NN	Pure ML atmospheric model
Lanthaler et al	2022	Neural Operator	PDE solver
Leufen & Schadler	2019	Fully connected NN	Parameterization
Li et al	2020	Graph NN	PDE solver
Li et al	2020	Neural Operator	PDE solver
Li et al	2020	Neural Operator	PDE solver
Lopez-Gomez et al	2023	Convolutional NN	Prediction
Lu et al	2020	Neural Operator	PDE solver
Meyer et al	2022	Fully connected NN	Emulation
Moishin et al	2021	Convolutional Recurrent NN	Prediction
Moore et al	2021	Fully connected NN	Emulation
Mudigonda et al	2017	Mixed/Custom NN	Object detection
Nelsen & Stuart	2021	Random Feature Model	PDE solver
Nguyen et al	2023	Mixed/Custom NN	Pure ML atmospheric model
O'Brien et al	2020	Bayesian model	Object detection
O'Gorman & Dwyer	2018	Decision tree-based	Emulation
O'Leary et al	2022	Fully connected NN	PDE solver
Ott et al	2020	Fully connected NN	Emulation
Pan et al	2020	Decision tree-based	Parameterisation
Patel et al	2021	Neural Operator	PDE solver
Pathak et al	2022	Mixed/Custom NN	Pure ML atmospheric model
Peairs & Chen	2011	Mixed/Custom non-NN	PDE solver
Pelissier et al	2020	Mixed/Custom non-NN	Hybrid model corrector
Prabhat et al	2021	Convolutional NN	Object detection
Psaros et al	2023	Neural Operator, Physics Informed NN	PDE solver
Rasp	2020	Fully connected NN	Emulation
Rasp et al	2018	Fully connected NN	Emulation
Rasp et al	2020	Fully connected NN, Linear regression	Pure ML atmospheric model
Rasp & Thuerey	2021	ResNet	Pure ML atmospheric model
Rizzuti et al	2019	Convolutional NN	NN based corrector step in PDE solver
Rosier et al	2023	Mixed/Custom NN	Prediction
Ross et al.	2023	Genetic programming, Linear regression, Convolutional NN	Intercomparison of methods to learn parameterisations from data

Rupe et al	2023	Mixed/Custom non-NN	Object detection
Sawada	2020	Regression	Emulation
Scher	2018	Convolutional NN	Emulation
Scher and Messori	2019	Convolutional NN	Emulation
Taylor & Feng	2022	Convolutional NN	Prediction
Tompson et al	2017	Convolutional NN	PDE solver
Toms et al	2020	Fully connected NN	NN interpretability
Ukkonen & Mäkelä	2019	Decision tree-based, Logistic Regression, Fully connected NN	Paramterisation
Ukkonen et al	2020	Fully connected NN	Emulation
Vlachas et al	2018	Recurrent NN	Pure ML baseline model
Wang et al	2021	Neural Operator	PDE solver
Wang et al	2022	ResNet	Parametrization
Wang et al	2022	Physics Informed NN	PDE solver
Watt-Meyer et al	2021	Decision tree-based	Nudging
Watson-Parris et al	2022	Gaussian Process, Decision tree-based, Mixed/Custom NN	Pure ML baseline model
Weyn et al	2019	Convolutional NN	Pure ML atmospheric model
Weyn et al	2020	Convolutional NN	Pure ML atmospheric model
Weyn et al	2021	Convolutional NN	Pure ML atmospheric model
Wikner et al	2020	Reservoir computing	Alongside-model bias corrector
Wu & Xiu	2020	ResNet	Learning PDE operators
Yamada et al	2018	Convolutional NN	Preconditioner
Yang et al	2016	Fully connected NN	PDE solver
Yeo et al	2021	Recurrent NN	Dynamical system simulation
Yuval & O’Gorman	2020	Decision tree-based	Emulation
Yuval et al	2021	Fully connected NN	Emulation
Zanna and Bolton	2020	Convolutional NN, Relevance vector machine	Parametrization and equation discovery
Zhao et al	2019	Fully connected NN	Paramterisation
Zhao et al	2019	Physics Informed NN	Paramterisation
Zhong et al	2023	Fully connected NN, Recurrent NN	Emulation

2396
2397

2398 **Code Availability**

2399 No code was used in the preparation of this review.

2400 **Data Availability**

2401 No data was processed in the preparation of this review except for the list of ML model types by cited paper, which
2402 is provided in the appendix.

Deleted: .

2403 **Author Contribution**

2404 COdBD researched and wrote Sections 3, 4, 5, 6 and 7, and provided review of sections 8, 10, and the glossary. TL
2405 researched and wrote sections 8, 10, and the glossary, and provided review of sections 3, 4, 5, 6, and 7. COdBD and
2406 TL researched and co-wrote sections 1, 2, 9, 11, 12, and the Appendix.

- Deleted: 2,
- Deleted: 8, and9,10
- Deleted: 11
- Deleted: 2,
- Deleted: and
- Deleted: 9
- Deleted: and
- Deleted: 0

2407 **Competing Interests**

2408 The authors declare that they have no conflict of interest.

2409 **Acknowledgements**

2410 The authors would like to thank Bethan White, Harrison Cook, Tom Dunstan and Karina Williams for their very
2411 helpful reviews of early versions of this manuscript. We also would like to wholeheartedly thank the referees for their
2412 extremely helpful, positive and well considered feedback and suggestions. Their input has greatly improved this
2413 review. Finally, we would like to acknowledge and thank the people who contacted us with comments, suggestions,
2414 and advice on the preprint versions of this review. All of the input was valuable, and greatly appreciated.

2415 **References**

Deleted: ¶

2416 Ackmann, J., Düben, P. D., Palmer, T. N., & Smolarkiewicz, P. K. Machine-learned preconditioners for linear solvers
2417 in geophysical fluid flows. *arXiv preprint arXiv:2010.02866*. <https://doi.org/10.48550/arXiv.2010.02866>. 6
2418 October 2020.

- Deleted: (2020).
- Formatted: Indent: Left: 0 cm, Hanging: 1.27 cm, Don't swap indents on facing pages

2419 Adams, S. V., Ford, R. W., Hambley, M., Hobson, J. M., Kavčić, I., Maynard, C. M., ... & Wong, R. LFRic: Meeting
2420 the challenges of scalability and performance portability in Weather and Climate models. *J PARALLEL*
2421 DISTR COM, 132, 383-396. <https://doi.org/10.1016/j.jpdc.2019.02.007>. 2019.

Deleted: Adams, S. V., Ford, R. W., Hambley, M., Hobson, J. M., Kavčić, I., Maynard, C. M., ... & Wong, R. (2019). LFRic: Meeting the challenges of scalability and performance portability in Weather and Climate models. *Journal of Parallel and Distributed Computing*, 132, 383-396.

2422 Alemohammad, S. H., Fang, B., Konings, A. G., Aires, F., Green, J. K., Kolassa, J., ... & Gentine, P. *Water, Energy,*
2423 and Carbon with Artificial Neural Networks (WECANN): a statistically based estimate of global surface
2424 turbulent fluxes and gross primary productivity using solar-induced fluorescence. *BIOGEOSCIENCES*,
2425 14(18), 4101-4124. <https://doi.org/10.5194/bg-14-4101-2017>. 2017.

- Deleted: . (2019)
- Deleted: . (2017).

2426 Andersson, T. R., Hosking, J. S., Pérez-Ortiz, M., Paige, B., Elliott, A., Russell, C., ... & Shuckburgh, E. *Seasonal*
2427 Arctic sea ice forecasting with probabilistic deep learning. *NAT COMMUN*, 12(1), 5124.
2428 <https://doi.org/10.1038/s41467-021-25257-4>. 2021.

Deleted: (2021).

2449 [Arcomano, T., Szunyogh, I., Wikner, A., Pathak, J., Hunt, B. R., & Ott, E. A Hybrid Approach to Atmospheric](#)
2450 [Modeling That Combines Machine Learning With a Physics-Based Numerical Model. J ADV MODEL](#)
2451 [EARTH SY, 14\(3\), e2021MS002712. <https://doi.org/10.1029/2021MS002712>. 2022.](#)
2452 Atkinson, S. Bayesian hidden physics models: Uncertainty quantification for discovery of nonlinear partial differential
2453 operators from data. *arXiv preprint arXiv:2006.04228*. <https://doi.org/10.48550/arXiv.2006.04228>. 7 June
2454 2020.
2455 Bar-Sinai, Y., Hoyer, S., Hickey, J., & Brenner, M. P. Learning data-driven discretizations for partial differential
2456 equations. *Proceedings of the National Academy of Sciences*, 116(31), 15344-15349.
2457 <https://doi.org/10.1073/pnas.1814058116>. 2019.
2458 [Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., ... & Pascanu, R.](#)
2459 [Relational inductive biases, deep learning, and graph networks. arXiv preprint arXiv:1806.01261.](#)
2460 <https://doi.org/10.48550/arXiv.1806.01261>. 4 June 2018.
2461 Beucler, T., Rasp, S., Pritchard, M., & Gentine, P. Achieving conservation of energy in neural network emulators for
2462 climate modeling. *arXiv preprint arXiv:1906.06622*. <https://doi.org/10.48550/arXiv.1906.06622>. 15 June
2463 2019.
2464 [Beucler, T., Pritchard, M., Rasp, S., Ott, J., Baldi, P., & Gentine, P. Enforcing analytic constraints in neural](#)
2465 [networks emulating physical systems. PHYS REV LETT, 126\(9\), 098302.](#)
2466 <https://doi.org/10.1103/PhysRevLett.126.098302>. 2021.
2467 Bhattacharya, K., Hosseini, B., Kovachki, N. B., & Stuart, A. M. Model reduction and neural networks for parametric
2468 PDEs. *arXiv preprint arXiv:2005.03180*. <https://doi.org/10.48550/arXiv.2005.03180>. 7 May 2020.
2469 Bi, K., Xie, L., Zhang, H., Chen, X., Gu, X., & Tian, Q. Pangu-Weather: A 3D High-Resolution Model for Fast and
2470 Accurate Global Weather Forecast. *arXiv preprint arXiv:2211.02556*.
2471 <https://doi.org/10.48550/arXiv.2211.02556>. 3 November 2022.
2472 [Bihlo, A., & Popovych, R. O. Physics-informed neural networks for the shallow-water equations on the sphere. J](#)
2473 [COMPUT PHYS, 456, 111024. <https://doi.org/10.1016/j.jcp.2022.111024>. 2022.](#)
2474 [Bolton, T., & Zanna, L. Applications of deep learning to ocean data inference and subgrid parameterization. J ADV](#)
2475 [MODEL EARTH SY, 11\(1\), 376-399. <https://doi.org/10.1029/2018MS001472>. 2019.](#)
2476 Breiman, L. (2001). Random forests. *Machine learning*, 45(1), 5-32.
2477 [Brenowitz, N. D., & Bretherton, C. S. Prognostic validation of a neural network unified physics parameterization.](#)
2478 [GEOPHYS RES LETT, 45\(12\), 6289-6298. <https://doi.org/10.1029/2018GL078510>. 2018.](#)
2479 [Brenowitz, N. D., & Bretherton, C. S. Spatially extended tests of a neural network parametrization trained by coarse-](#)
2480 [graining. J ADV MODEL EARTH SY, 11\(8\), 2728-2744. <https://doi.org/10.1029/2019MS001711>. 2019.](#)
2481 [Brenowitz, N. D., Beucler, T., Pritchard, M., & Bretherton, C. S. Interpreting and stabilizing machine-learning](#)
2482 [parametrizations of convection. J ATMOS SCI, 77\(12\), 4357-4375. \[0082.1. 2020.\]\(https://doi.org/10.1175/JAS-D-20-

2483 <a href=\)](#)

Moved down [4]: (2022).

Deleted: Arcomano, T., Szunyogh, I., Wikner, A., Pathak, J., Hunt, B. R., & Ott, E. (2022). A Hybrid Approach to Atmospheric Modeling That Combines Machine Learning With a Physics-Based Numerical Model. *Journal of Advances in Modeling Earth Systems*, 14(3), e2021MS002712.

Moved (insertion) [4]

Deleted: (

Deleted:)

Deleted: (2020).

Formatted: Indent: Left: 0 cm, Hanging: 1.27 cm, Don't swap indents on facing pages

Field Code Changed

Deleted: . (2019)

Formatted: Indent: Left: 0 cm, Hanging: 1.27 cm, Don't swap indents on facing pages

Deleted: . (2018)

Deleted: . (2019)

Formatted: Indent: Left: 0 cm, Hanging: 1.27 cm, Don't swap indents on facing pages

Deleted: Beucler, T., Pritchard, M., Rasp, S., Ott, J., Baldi, P., & Gentine, P. (2021). Enforcing analytic constraints in neural networks emulating physical systems. *Physical Review Letters*, 126(9), 098302. [19]

Deleted: . (2021)

Formatted: Left, Indent: Left: 0 cm, Hanging: 1.27 cm

Field Code Changed

Deleted: (2020).

Formatted: (2022) [20]

Deleted: . (2022)

Formatted: (2022) [21]

Formatted: Font: Not Italic

Formatted: Font: Not Italic

Formatted: Font: Not Italic

Formatted: Font: Not Italic

Deleted: Bihlo, A., & Popovych, R. O. (2022). Physics-informed neural networks for the shallow-water equations on the sphere. *Journal of Computational Physics*, 456, 111024. [22]

Deleted: (2022).

Deleted: Bolton, T., & Zanna, L. (2019). Applications of deep learning to ocean data inference and subgrid parameterization. *Journal of Advances in Modeling Earth Systems*, 11(1), 376-399. [23]

Deleted: . (2019).

Deleted: Brenowitz, N. D., & Bretherton, C. S. (2018). Prognostic validation of a neural network unified physics parameterization. *Geophysical Research Letters*, 45(12), 6289-6298. [24]

Deleted: . (2018).

Deleted: Brenowitz, N. D., & Bretherton, C. S. (2019). Spatially extended tests of a neural network parametrization trained by coarse-graining. *Journal of Advances in Modeling Earth Systems*, 11(8), 2728-2744. [25]

Deleted: (2019).

Deleted: Brenowitz, N. D., Beucler, T., Pritchard, M., & Bretherton, C. S. (2020). Interpreting and stabilizing machine-learning parametrizations of convection. *Journal of Atmospheric Sciences*, 77(12), 4357-4375. [26]

Deleted: (2020).

2536 Brenowitz, N. D., Henn, B., McGibbon, J., Clark, S. K., Kwa, A., Perkins, W. A., ... & Bretherton, C. S. Machine
 2537 learning climate model dynamics: Offline versus online performance. *arXiv preprint arXiv:2011.03081*.
 2538 <https://doi.org/10.48550/arXiv.2011.03081>. 5 November 2020.

2539 Brenowitz, N. D., Perkins, W. A., Nugent, J. M., Watt-Meyer, O., Clark, S. K., Kwa, A., ... & Bretherton, C. S.
 2540 Emulating Fast Processes in Climate Models. *arXiv preprint arXiv:2211.10774*.
 2541 <https://doi.org/10.48550/arXiv.2211.10774>. 19 November 2022.

2542 Carranza-García, M., García-Gutiérrez, J., & Riquelme, J. C. A framework for evaluating land use and land cover
 2543 classification using convolutional neural networks. *REMOTE SENS-BASEL*, 11(3), 274.
 2544 <https://doi.org/10.3390/rs11030274>. 2019.

2545 Chaney, N. W., Herman, J. D., Ek, M. B., & Wood, E. F. Deriving global parameter estimates for the Noah land
 2546 surface model using FLUXNET and machine learning. *J GEOPHYS RES-ATMOS*, 121(22), 13-218.
 2547 <https://doi.org/10.1002/2016JD024821>. 2016.

2548 Chantry, M., Christensen, H., Dueben, P., & Palmer, T. Opportunities and challenges for machine learning in weather
 2549 and climate modelling: hard, medium and soft AI. *Philosophical Transactions of the Royal Society A*,
 2550 379(2194), 20200083. <https://doi.org/10.1098/rsta.2020.0083>. 2021.

2551 Chantry, M., Hatfield, S., Dueben, P., Polichtchouk, I., & Palmer, T. Machine learning emulation of gravity wave
 2552 drag in numerical weather forecasting. *J ADV MODEL EARTH SY*, 13(7), e2021MS002477.
 2553 <https://doi.org/10.1029/2021MS002477>. 2021.

2554 Chase, R. J., Harrison, D. R., Burke, A., Lackmann, G. M., & McGovern, A. (2022a). A Machine Learning Tutorial
 2555 for Operational Meteorology, Part I: Traditional Machine Learning. *arXiv preprint arXiv:2204.07492*.
 2556 <https://doi.org/10.48550/arXiv.2204.07492>. 15 April 2022.

2557 Chase, R. J., Harrison, D. R., Lackmann, G., & McGovern, A. A Machine Learning Tutorial for Operational
 2558 Meteorology, Part II: Neural Networks and Deep Learning. *arXiv preprint arXiv:2211.00147*.
 2559 <https://doi.org/10.48550/arXiv.2211.00147>. 31 October 2022.

2560 Chattopadhyay, A., Subel, A., & Hassanzadeh, P. Data-driven super-parameterization using deep learning:
 2561 Experimentation with multiscale Lorenz 96 systems and transfer learning. *J ADV MODEL EARTH SY*,
 2562 12(11), e2020MS002084. <https://doi.org/10.1029/2020MS002084>. 2020.

2563 Chevallier, F., Chérut, F., Scott, N. A., & Chédin, A. A neural network approach for a fast and accurate computation
 2564 of a longwave radiative budget. *J APPL METEOROL*, 37(11), 1385-1397. [https://doi.org/10.1175/1520-0450\(1998\)037%3C1385:ANNAFA%3E2.0.CO;2](https://doi.org/10.1175/1520-0450(1998)037%3C1385:ANNAFA%3E2.0.CO;2). 1998.

2565 Chi, J., & Kim, H. C. Prediction of arctic sea ice concentration using a fully data driven deep neural network. *REMOTE*
 2566 *SENS-BASEL*, 9(12), 1305. <https://doi.org/10.3390/rs9121305>. 2017.

2567 Clare, M. C., Jamil, O., & Morcrette, C. J. (2021). Combining distribution-based neural networks to predict weather
 2568 forecast probabilities. *Q J ROY METEOR SOC*, 147(741), 4337-4357. <https://doi.org/10.1002/qj.4180>.
 2569 2021.

2570 Conley, A. J., Garcia, R., Kinnison, D., Lamarque, J. F., Marsh, D., Mills, M., ... & Taylor, M. A. Description of the
 2571 NCAR community atmosphere model (CAM 5.0). NCAR technical note, 3. 2012.

Deleted: (2020).

Deleted: (2022).

Formatted: Font: Not Italic

Formatted: Font: Not Italic

Formatted: Font: Not Italic

Formatted: Font: Not Italic

Deleted: (2019).

Formatted: Font: Not Italic

Formatted: Font: Not Italic

Formatted: Font: Not Italic

Deleted: (2016).

Formatted: Font: Not Italic

Deleted: (2021).

Deleted: Chantry, M., Hatfield, S., Dueben, P., Polichtchouk, I., & Palmer, T. (2021). Machine learning emulation of gravity wave drag in numerical weather forecasting. *Journal of Advances in Modeling Earth Systems*, 13(7), e2021MS002477.

Deleted: (2021).

Deleted: .

Formatted: Font: Not Italic

Formatted: Font: Not Italic

Deleted: . (2022b)

Formatted: Font: Not Italic

Formatted: Font: Not Italic

Formatted: Font: Not Italic

Formatted: Font: Not Italic

Deleted: Chattopadhyay, A., Subel, A., & Hassanzadeh, P. (2020). Data-driven super-parameterization using deep learning: Experimentation with multiscale Lorenz 96 systems and transfer learning. *Journal of Advances in Modeling Earth Systems*, 12(11), e2020MS002084.

Deleted: (2020).

Deleted: Chevallier, F., Chérut, F., Scott, N. A., & Chédin, A. (1998). A neural network approach for a fast and accurate computation of a longwave radiative budget. *Journal of applied meteorology*, 37(11), 1385-1397.

Deleted: . (1998).

Deleted: (2017).

Deleted: Clare, M. C., Jamil, O., & Morcrette, C. J. (2021). Combining distribution-based neural networks to predict weather forecast probabilities. *Quarterly Journal of the Royal Meteorological Society*, 147(741), 4337-4357.

Deleted: (2012).

2603 Cuomo, S., Di Cola, V. S., Giampaolo, F., Rozza, G., Raissi, M., & Piccialli, F. Scientific Machine Learning through
2604 Physics-Informed Neural Networks: Where we are and What's next. *arXiv preprint arXiv:2201.05624*.
2605 <https://doi.org/10.48550/arXiv.2201.05624>. 14 January 2022. Deleted: (2022).

2606 Dagon, K., Sanderson, B. M., Fisher, R. A., & Lawrence, D. M. A machine learning approach to emulation and
2607 biophysical parameter estimation with the Community Land Model, version 5. *Advances in Statistical*
2608 *Climatology, Meteorology and Oceanography*, 6(2), 223-244. <https://doi.org/10.5194/ascmo-6-223-2020>.
2609 2020. Deleted: (2020).

2610 De Bézenac, E., Pajot, A., & Gallinari, P. *Towards a hybrid approach to physical process modeling*. Technical report.
2611 2017. Deleted: . (2017)

2612 de Witt, C. S., Tong, C., Zantedeschi, V., De Martini, D., Kalaitzis, F., Chantry, M., ... & Bilinski, P. *RainBench:*
2613 *towards global precipitation forecasting from satellite imagery*. *arXiv preprint arXiv:2012.09670*.
2614 <https://doi.org/10.48550/arXiv.2012.09670>. 17 December 2020. Deleted: (2020).

2615 Deng, J., Dong, W., Socher, R., Li, L. J., Li, K., & Fei-Fei, L. *Imagenet: A large-scale hierarchical image database*.
2616 *PROC CVPR IEEE* (pp. 248-255). Ieee. <https://doi.org/10.1109/CVPR.2009.5206848>. 2009. Deleted: Deng, J., Dong, W., Socher, R., Li, L. J., Li, K., & Fei-Fei, L. (2009, June). Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition* (pp. 248-255). Ieee.

2617 Digra, M., Dhir, R., & Sharma, N. Land use land cover classification of remote sensing images based on the deep
2618 learning approaches: a statistical analysis and review. *ARAB J GEOSCI*, 15(10), 1003.
2619 <https://doi.org/10.1007/s12517-022-10246-8>. 2022. Deleted: (2009, June).

2620 Dijkstra, H. A., Petersik, P., Hernández-García, E., & López, C. The application of machine learning techniques to
2621 improve El Niño prediction skill. *AIP CONF PROC*, 153. <https://doi.org/10.3389/fphy.2019.00153>. 2019. Deleted: In 2009

2622 Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... & Houlsby, N. An image is
2623 worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
2624 <https://doi.org/10.48550/arXiv.2010.11929>. 22 October 2020. Deleted: . (2022).

2625 Dueben, P. D., & Bauer, P. Challenges and design choices for global weather and climate models based on machine
2626 learning. *GEOSCI MODEL DEV*, 11(10), 3999-4009. <https://doi.org/10.5194/gmd-11-3999-2018>. 2018. Deleted: (2019).

2627 Dueben, P. D., Schultz, M. G., Chantry, M., Gagne, D. J., Hall, D. M., & McGovern, A. Challenges and Benchmark
2628 Datasets for Machine Learning in the Atmospheric Sciences: Definition, Status, and Outlook. *Artificial*
2629 *Intelligence for the Earth Systems*, 1(3), e210002. <https://doi.org/10.1175/AIES-D-21-0002.1>. 2022. Deleted: *Frontiers in Physics*

2630 ECMWF. (2018). Ifs documentation (cy45r1). Retrieved from [https://www.ecmwf.int/en/publications/ifs-](https://www.ecmwf.int/en/publications/ifs-documentation)
2631 [documentation](https://www.ecmwf.int/en/publications/ifs-documentation) accessed 7th February 2023 Deleted: *Frontiers in Physics*

2632 Eyring, V., Bony, S., Meehl, G. A., Senior, C. A., Stevens, B., Stouffer, R. J., & Taylor, K. E. Overview of the Coupled
2633 Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization. *GEOSCI MODEL*
2634 *DEV*, 9(5), 1937-1958. <https://doi.org/10.5194/gmd-9-1937-2016>. 2016. Deleted: (2020).

2635 Flora, M., Potvin, C., McGovern, A., & Handler, S. Comparing Explanation Methods for Traditional Machine
2636 Learning Models Part 2: Quantifying Model Explainability Faithfulness and Improvements with
2637 Dimensionality Reduction. *arXiv preprint arXiv:2211.10378*. <https://doi.org/10.48550/arXiv.2211.10378>. 18
2638 November 2022. Deleted: (2018).

2639 Friedman, J. H. Greedy function approximation: a gradient boosting machine. *ANN STAT*, 1189-1232. 2001. Deleted: (2022).

2679 [Fuhg, J. N., Kalogeris, I., Fau, A., & Bouklas, N. Interval and fuzzy physics-informed neural networks for uncertain](#)
2680 [fields. *PROBABILIST ENG MECH.* 68, 103240. 2022.](#)

2681 Gagne, D. J., McCandless, T., Kosovic, B., DeCastro, A., Loft, R., Haupt, S. E., & Yang, B. Machine learning
2682 parameterization of the surface layer: bridging the observation-modeling gap. In *AGU Fall Meeting Abstracts*
2683 (Vol. 2019, pp. IN44A-04). 2019.

2684 [Gagne, D. J., Chen, C. C., & Gettelman, A. Emulation of bin Microphysical Processes with machine learning. In *100th*](#)
2685 [American Meteorological Society Annual Meeting. AMS. 2020.](#)

2686 Gagne, D. J., Christensen, H. M., Subramanian, A. C., & Monahan, A. H. Machine learning for stochastic
2687 parameterization: Generative adversarial networks in the Lorenz'96 model. *J ADV MODEL EARTH SY.*
2688 *12*(3), e2019MS001896. <https://doi.org/10.1029/2019MS001896>. 2020.

2689 [Garg, S., Rasp, S., & Thuerey, N. WeatherBench Probability: A benchmark dataset for probabilistic medium-range](#)
2690 [weather forecasting along with deep learning baseline models. arXiv preprint arXiv:2205.00865.](#)
2691 <https://doi.org/10.48550/arXiv.2205.00865>. 2 May 2022.

2692 [George, T., Gupta, A., & Sarin, V. A recommendation system for preconditioned iterative solvers. *IEEE DATA*](#)
2693 [MINING \(pp. 803-808\). IEEE. <https://doi.org/10.1109/ICDM.2008.105>. 2008.](#)

2694 [Gettelman, A., Gagne, D. J., Chen, C. C., Christensen, M. W., Lebo, Z. J., Morrison, H., & Gantos, G. Machine](#)
2695 [learning the warm rain process. *J ADV MODEL EARTH SY.* 13\(2\), e2020MS002268.](#)
2696 <https://doi.org/10.1029/2020MS002268>. 2021.

2697 [Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. Generative](#)
2698 [adversarial networks. *COMMUN ACM.* 63\(11\), 139-144. <https://doi.org/10.1145/3422622>. 2020.](#)

2699 Goodfellow, I., Yoshua B., & Aaron C. *Deep learning. MIT press. 2016.*

2700 [Grigo, C., & Koutsourelakis, P. S. \(2019\). A physics-aware, probabilistic machine learning framework for coarse-](#)
2701 [graining high-dimensional systems in the Small Data regime. *J COMPUT PHYS.* 397, 108842.](#)
2702 <https://doi.org/10.1016/j.jcp.2019.05.053>. 2019.

2703 [Gurvan, M., Bourdallé-Badie, R., Chanut, J., Clementi, E., Coward, A., Ethé, C., ... & Samson, G. NEMO ocean](#)
2704 [engine. Institut Pierre-Simon Laplace \(IPSL\), Zenodo. 2019.](#)

2705 Ham, Y. G., Kim, J. H., & Luo, J. J. Deep learning for multi-year ENSO forecasts. *Nature*, 573(7775), 568-572.
2706 <https://doi.org/10.1038/s41586-019-1559-7>. 2019.

2707 [Ham, Y. G., Kim, J. H., Kim, E. S., & On, K. W. Unified deep learning model for El Niño/Southern Oscillation](#)
2708 [forecasts by incorporating seasonality in climate data. *SCI BULL.* 66\(13\), 1358-1366.](#)
2709 <https://doi.org/10.1016/j.scib.2021.03.009>. 2021.

2710 [Han, Y., Zhang, G. J., Huang, X., & Wang, Y. A moist physics parameterization based on deep learning. *J ADV*](#)
2711 [MODEL EARTH SY. 12\(9\), e2020MS002076. <https://doi.org/10.1029/2020MS002076>. 2020.](#)

2712 Harder, P., Watson-Parris, D., Stier, P., Strassel, D., Gauger, N. R., & Keuper, J. Physics-informed learning of aerosol
2713 microphysics. *Environmental Data Science*, 1, e20. <https://doi.org/10.1017/eds.2022.22>. 2022.

2714 Harris, L., Chen, X., Putman, W., Zhou, L., & Chen, J. H. A scientific description of the GFDL finite-volume cubed-
2715 sphere dynamical core. <https://doi.org/10.25923/6nhs-5897>. 2021.

Deleted: Fuhg, J. N., Kalogeris, I., Fau, A., & Bouklas, N. (2022). Interval and fuzzy physics-informed neural networks for uncertain fields. *Probabilistic Engineering Mechanics*, 68, 103240.

Deleted: (2022).

Deleted: (2019, December)

Deleted: (2020, January).

Deleted: (2020, February).

Deleted: *Journal of Advances in Modeling Earth Systems*

Deleted: Gagne, D. J., Chen, C. C., & Gettelman, A. (2020, January). Emulation of bin Microphysical Processes with machine learning. In *100th American Meteorological Society Annual Meeting*. AMS.

Deleted: (2022).

Deleted: George, T., Gupta, A., & Sarin, V. (2008, December). A recommendation system for preconditioned iterative solvers. In *2008 Eighth IEEE International Conference on Data Mining (pp. 803-808)*. IEEE.

Deleted: . (2008, December)

Deleted: In 2008 Eighth

Deleted: Gettelman, A., Gagne, D. J., Chen, C. C., Christensen, M. W., Lebo, Z. J., Morrison, H., & Gantos, G. (2021). Machine learning the warm rain process. *Journal of Advances in Modeling Earth Systems*, 13(2), e2020MS002268.

Deleted: (2021).

Deleted: Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. (2020). Generative adversarial networks. *Communications of the ACM*, 63(11), 139-144.

Deleted: (2020).

Deleted: (2016).

Deleted: Grigo, C., & Koutsourelakis, P. S. (2019). A physics-aware, probabilistic machine learning framework for coarse-graining high-dimensional systems in the Small Data regime. *Journal of Computational Physics*, 397, 108842.

Deleted: ...

Deleted: . (2019)

Deleted: . (2019).

Deleted: Ham, Y. G., Kim, J. H., Kim, E. S., & On, K. W. (2021). Unified deep learning model for El Niño/Sout... [27]

Deleted: (2021).

Deleted: Han, Y., Zhang, G. J., Huang, X., & Wang, Y. [28]

Deleted: (2020).

Deleted: (2022).

Deleted: (2021).

2855 Keisler, R. Forecasting Global Weather with Graph Neural Networks. *arXiv preprint arXiv:2202.07575*.
2856 <https://doi.org/10.48550/arXiv.2202.07575>. 15 February 2022. Deleted: ¶

2857 Kelotra, A., & Pandey, P. (2020). Stock market prediction using optimized deep-convlstm model. *Big Data*, 8(1), 5-
2858 24. <https://doi.org/10.48550/arXiv.2202.07575>. 11 February 2022. Deleted: ¶

2859 Kim, J., Kwon, M., Kim, S. D., Kug, J. S., Ryu, J. G., & Kim, J. Spatiotemporal neural network with attention
2860 mechanism for El Niño forecasts. *SCI REP-UK*, 12(1), 1-15. <https://doi.org/10.1038/s41598-022-10839-z>. Deleted: Kim, J., Kwon, M., Kim, S. D., Kug, J. S., Ryu, J. G., & Kim, J. (2022). Spatiotemporal neural network with
2861 2022. attention mechanism for El Niño forecasts. *Scientific Reports*, 12(1), 1-15. Deleted: (2022).

2862 Kipf, T. N., & Welling, M. Semi-supervised classification with graph convolutional networks. *arXiv preprint*
2863 *arXiv:1609.02907*. <https://doi.org/10.48550/arXiv.1609.02907>. 9 September 2016. Deleted: (2022). Deleted: (2016).

2864 Kochkov, D., Smith, J. A., Alieva, A., Wang, Q., Brenner, M. P., & Hoyer, S. Machine learning-accelerated
2865 computational fluid dynamics. *Proceedings of the National Academy of Sciences*, 118(21), e2101784118.
2866 <https://doi.org/10.1073/pnas.2101784118>. 2021. Deleted: (2021).

2867 Krasnopolsky, V. M., Chalikov, D. V., & Tolman, H. L. A neural network technique to improve computational
2868 efficiency of numerical oceanic models. *OCEAN MODEL*, 4(3-4), 363-383. [https://doi.org/10.1016/S1463-5003\(02\)00010-0](https://doi.org/10.1016/S1463-5003(02)00010-0). 2002. Deleted: Krasnopolsky, V. M., Chalikov, D. V., & Tolman, H. L. (2002). A neural network technique to improve computational efficiency of numerical oceanic models. *Ocean Modelling*, 4(3-4), 363-383. Deleted: (2002).

2870 Krasnopolsky, V. M., Fox-Rabinovitz, M. S., & Chalikov, D. V. New approach to calculation of atmospheric model
2871 physics: Accurate and fast neural network emulation of longwave radiation in a climate model. *MON*
2872 *WEATHER REV*, 133(5), 1370-1383. <https://doi.org/10.1175/MWR2923.1>. 2005. Deleted: Krasnopolsky, V. M., Fox-Rabinovitz, M. S., & Chalikov, D. V. (2005). New approach to calculation of atmospheric model physics: Accurate and fast neural network emulation of longwave radiation in a climate model. *Monthly Weather Review*, 133(5), 1370-1383. Deleted: (2005).

2873 Krasnopolsky, V. M., Fox-Rabinovitz, M. S., & Belochitski, A. A. Using ensemble of neural networks to learn
2874 stochastic convection parameterizations for climate and numerical weather prediction models from data
2875 simulated by a cloud resolving model. *Advances in Artificial Neural Systems*, 2013. Deleted: (2013).
2876 <https://doi.org/10.1155/2013/485913>. 2013. Deleted: (2013).

2877 Kuefler, E., & Chen, T. Y. On using reinforcement learning to solve sparse linear systems. In *International Conference*
2878 *on Computational Science* (pp. 955-964). Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-540-69384-0_100. 2008. Deleted: (2008, June).

2880 Ladický, L. U., Jeong, S., Solenthaler, B., Pollefeys, M., & Gross, M. Data-driven fluid simulations using regression
2881 forests. *ACM T GRAPHIC (TOG)*, 34(6), 1-9. <https://doi.org/10.1145/2816795.2818129>. 2015. Deleted: Ladický, L. U., Jeong, S., Solenthaler, B., Pollefeys, M., & Gross, M. (2015). Data-driven fluid simulations using regression forests. *ACM Transactions on Graphics (TOG)*, 34(6), 1-9. Deleted: (2015).

2882 Lam, R., Sanchez-Gonzalez, A., Willson, M., Wimsberger, P., Fortunato, M., Pritzel, A., ... & Battaglia, P. GraphCast:
2883 Learning skillful medium-range global weather forecasting. *arXiv preprint arXiv:2212.12794*. Deleted: (2022).
2884 <https://doi.org/10.48550/arXiv.2212.12794>. 24 December 2022. Deleted: (2022).

2885 Lanthaler, S., Mishra, S., & Karniadakis, G. E. Error estimates for deepoanets: A deep learning framework in infinite
2886 dimensions. *Transactions of Mathematics and Its Applications*, 6(1), tnac001.
2887 <https://doi.org/10.1093/imatrn/tnac001>. 2022. Deleted: (2022).

2888 LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. Gradient-based learning applied to document recognition. *P IEEE*,
2889 86(11), 2278-2324. <https://doi.org/10.1109/5.726791>. 1998. Deleted: LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 2278-2324. Deleted: (1998).

2890 Leufen, L. H., & Schädlér, G. Calculating the turbulent fluxes in the atmospheric surface layer with neural networks.
2891 *GEOSCI MODEL DEV*, 12(5), 2033-2047. <https://doi.org/10.5194/gmd-12-2033-2019>. 2019. Deleted: (2019).

2927 [Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Stuart, A., Bhattacharya, K., & Anandkumar, A. Multipole graph](#)
2928 [neural operator for parametric partial differential equations. *ADV NEUR IN*, 33, 6755-6766. 2020a.](#)
2929 Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., & Anandkumar, A. Neural operator:
2930 Graph kernel network for partial differential equations. *arXiv preprint arXiv:2003.03485*.
2931 <https://doi.org/10.48550/arXiv.2003.03485>. 7 March 2020.
2932 Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., & Anandkumar, A. (2020). Fourier
2933 neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*.
2934 <https://doi.org/10.48550/arXiv.2010.08895>. 18 October 2020.
2935 [Lopez-Gomez, I., McGovern, A., Agrawal, S., & Hickey, J. Global extreme heat forecasting using neural weather](#)
2936 [models. *Artificial Intelligence for the Earth Systems*, 2\(1\), e220035. https://doi.org/10.1175/AIES-D-22-](#)
2937 [0035.1. 2023.](#)
2938 Lorenz, E. N. Predictability: A problem partly solved, in Proceedings of Seminar on Predictability, 4–8 September
2939 1995. <https://doi.org/10.1017/CBO9780511617652.004>. 1995
2940 Lu, L., Jin, P., & Karniadakis, G. E. Deeponet: Learning nonlinear operators for identifying differential equations
2941 based on the universal approximation theorem of operators. *arXiv preprint arXiv:1910.03193*.
2942 <https://doi.org/10.48550/arXiv.1910.03193>. 8 October 2019.
2943 [Lundberg S., & Lee S. A Unified Approach to Interpreting Model Predictions. *ADV NEUR IN*, 30, 4768–4777. 2017.](#)
2944 [McCulloch, W. S., & Pitts, W. A logical calculus of the ideas immanent in nervous activity. *The B MATH BIOPHYS*,](#)
2945 [5\(4\), 115-133. https://doi.org/10.1007/BF02478259. 1943.](#)
2946 [McGovern, A., Lagerquist, R., Gagne, D. J., Jergensen, G. E., Elmore, K. L., Homeyer, C. R., & Smith, T. Making](#)
2947 [the black box more transparent: Understanding the physical implications of machine learning. *B AM*](#)
2948 [METEOROL SOC](#), 100(11), 2175-2199. <https://doi.org/10.1175/BAMS-D-18-0195.1>. 2019.
2949 [Meyer, D., Hogan, R. J., Dueben, P. D., & Mason, S. L. Machine learning emulation of 3D cloud radiative effects. *J*](#)
2950 [ADV MODEL EARTH SY](#), 14(3), e2021MS002550. <https://doi.org/10.1029/2021MS002550>. 2022.
2951 [Moishin, M., Deo, R. C., Prasad, R., Raj, N., & Abdulla, S. Designing deep-based learning flood forecast model with](#)
2952 [ConvLSTM hybrid algorithm. *IEEE ACCESS*, 9, 50982-50993.](#)
2953 <https://doi.org/10.1109/ACCESS.2021.3065939>. 2021.
2954 [Molina, M. J., O'Brien, T. A., Anderson, G., Ashfaq, M., Bennett, K. E., Collins, W. D., ... & Ullrich, P. A. A Review](#)
2955 [of Recent and Emerging Machine Learning Applications for Climate Variability and Weather Phenomena.](#)
2956 [Artificial Intelligence for the Earth Systems](#), 1-46. <https://doi.org/10.1175/AIES-D-22-0086.1>. 2023.
2957 [Mooers, G., Pritchard, M., Beucler, T., Ott, J., Yacalis, G., Baldi, P., & Gentine, P. \(2021\). Assessing the Potential of](#)
2958 [Deep Learning for Emulating Cloud Superparameterization in Climate Models With Real-Geography](#)
2959 [Boundary Conditions. *J ADV MODEL EARTH SY*, 13\(5\), e2020MS002385.](#)
2960 <https://doi.org/10.1029/2020MS002385>. 2021.
2961 Mudigonda, M., Kim, S., Mahesh, A., Kahou, S., Kashinath, K., Williams, D., ... & Prabhat, M. Segmenting and
2962 tracking extreme climate events using neural networks. In *Deep Learning for Physical Sciences (DLPS)*
2963 *Workshop, held with NIPS Conference. 2017.*

Deleted: Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Stuart, A., Bhattacharya, K., & Anandkumar, A. (2020a). Multipole graph neural operator for parametric partial differential equations. *Advances in Neural Information Processing Systems*, 33, 6755-6766.

Deleted: (2020a).

Deleted: (2020b).

Deleted: c

Deleted: (2023).

Deleted: (1996).

Deleted: (2019).

Deleted: ¶

Deleted: (2017).

Deleted: McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5(4), 115-133.

Deleted: (1943).

Deleted: McGovern, A., Lagerquist, R., Gagne, D. J., Jergensen, G. E., Elmore, K. L., Homeyer, C. R., & Smith, T. (2019). Making the black box more transparent: Understanding the physical implications of machine learning. *Bulletin of the American Meteorological Society*, 100(11), 2175-2199.

Deleted: (2019).

Deleted: Meyer, D., Hogan, R. J., Dueben, P. D., & Mason, S. L. (2022). Machine learning emulation of 3D cloud radiative effects. *Journal of Advances in Modeling Earth Systems*, 14(3), e2021MS002550.

Deleted: (2022).

Deleted: Moishin, M., Deo, R. C., Prasad, R., Raj, N., & Abdulla, S. (2021). Designing deep-based learning flood forecast model with ConvLSTM hybrid algorithm. *IEEE Access*, 9, 50982-50993.

Deleted: (2021).

Deleted: (2023).

Deleted: Mooers, G., Pritchard, M., Beucler, T., Ott, J., Yacalis, G., Baldi, P., & Gentine, P. (2021). Assessing the Potential of Deep Learning for Emulating Cloud Superparameterization in Climate Models With Real-Geography Boundary Conditions. *Journal of Advances in Modeling Earth Systems*, 13(5), e2020MS002385.

Deleted: (2017).

3006 [Nelsen, N. H., & Stuart, A. M. The random feature model for input-output maps between banach spaces. SIAM J SCI](#)
3007 [COMPUT, 43\(5\), A3212-A3243. <https://doi.org/10.1137/20M133957X>. 2021.](#)

3008 [Nguyen, T., Brandstetter, J., Kapoor, A., Gupta, J. K., & Grover, A. ClimaX: A foundation model for weather and](#)
3009 [climate. arXiv preprint arXiv:2301.10343. <https://doi.org/10.48550/arXiv.2301.10343>. 24 January 2023.](#)

3010 [Nielsen, A. H., Iosifidis, A., & Karstoft, H. Forecasting large-scale circulation regimes using deformable](#)
3011 [convolutional neural networks and global spatiotemporal climate data. SCI REP-UK, 12\(1\), 1-12.](#)
3012 [\[https://doi.org/10.1175/1520-0469\\(1995\\)052%3C1237:WRAQS%3E2.0.CO;2\]\(https://doi.org/10.1175/1520-0469\(1995\)052%3C1237:WRAQS%3E2.0.CO;2\). 2022.](#)

3013 [O'Brien, T. A., Risser, M. D., Loring, B., Elbashandy, A. A., Krishnan, H., Johnson, J., ... & Collins, W. D. Detection](#)
3014 [of atmospheric rivers with inline uncertainty quantification: TECA-BARD v1. 0.1. GEOSCI MODEL DEV,](#)
3015 [13\(12\), 6131-6148. <https://doi.org/10.5194/gmd-13-6131-2020>. 2020.](#)

3016 [O'Gorman, P. A., & Dwyer, J. G. Using machine learning to parameterize moist convection: Potential for modeling](#)
3017 [of climate, climate change, and extreme events. J ADV MODEL EARTH SY, 10\(10\), 2548-2563.](#)
3018 [<https://doi.org/10.1029/2018MS001351>. 2018.](#)

3019 [O'Leary, J., Paulson, J. A., & Mesbah, A. Stochastic physics-informed neural ordinary differential equations. J](#)
3020 [COMPUT PHYS, 468, 111466. <https://doi.org/10.1016/j.jcp.2022.111466>. 2022.](#)

3021 [Ott, J., Pritchard, M., Best, N., Linstead, E., Curcic, M., & Baldi, P. A Fortran-Keras deep learning bridge for scientific](#)
3022 [computing. Scientific Programming, 2020. <https://doi.org/10.1155/2020/8888811>. 2020.](#)

3023 [Pal, S., & Sharma, P. A review of machine learning applications in land surface modeling. Earth, 2\(1\), 174-190.](#)
3024 [<https://doi.org/10.3390/earth2010011>. 2021.](#)

3025 [Palmer, T. A vision for numerical weather prediction in 2030. arXiv preprint arXiv:2007.04830.](#)
3026 [<https://doi.org/10.48550/arXiv.2007.04830>. 3 July 2020.](#)

3027 [Pan, S., Pan, N., Tian, H., Friedlingstein, P., Sitch, S., Shi, H., ... & Running, S. W. Evaluation of global terrestrial](#)
3028 [evapotranspiration using state-of-the-art approaches in remote sensing, machine learning and land surface](#)
3029 [modeling. HYDROL EARTH SYST SC, 24\(3\), 1485-1509. <https://doi.org/10.5194/hess-24-1485-2020>.](#)
3030 [2020.](#)

3031 [Patel, R. G., Trask, N. A., Wood, M. A., & Cyr, E. C. A physics-informed operator regression framework for extracting](#)
3032 [data-driven continuum models. COMPUT METHOD APPL M, 373, 113500.](#)
3033 [<https://doi.org/10.1016/j.cma.2020.113500>. 2021.](#)

3034 [Pathak, J., Subramanian, S., Harrington, P., Raja, S., Chattopadhyay, A., Mardani, M., ... & Anandkumar, A. \(2022\).](#)
3035 [Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators.](#)
3036 [arXiv preprint arXiv:2202.11214. <https://doi.org/10.48550/arXiv.2202.11214>. 22 February 2022.](#)

3037 [Peairs, L., & Chen, T. Y. Using reinforcement learning to vary the m in GMRES \(m\). PROCEDIA COMPUT SCI, 4,](#)
3038 [2257-2266. <https://doi.org/10.1016/j.procs.2011.04.246>. 2011.](#)

3039 [Pelissier, C., Frame, J., & Nearing, G. Combining parametric land surface models with machine learning. INT](#)
3040 [GEOSCI REMOTE SE, \(pp. 3668-3671\). IEEE. <https://doi.org/10.1109/IGARSS39084.2020.9324607>.](#)
3041 [2020.](#)

Deleted: Nelsen, N. H., & Stuart, A. M. (2021). The random feature model for input-output maps between banach spaces. *SIAM Journal on Scientific Computing*, 43(5), A3212-A3243.

Deleted: (2021).

Deleted: Nguyen, T., Brandstetter, J., Kapoor, A., Gupta, J. K., & Grover, A. (2023). ClimaX: A foundation model for weather and climate. *arXiv preprint arXiv:2301.10343*. Nielsen, A. H., Iosifidis, A., & Karstoft, H. (2022). Forecasting large-scale circulation regimes using deformable convolutional neural networks and global spatiotemporal climate data. *Scientific Reports*, 12(1), 1-12.

Deleted: (2023).

Deleted: (2022).

Formatted: Indent: Left: 0 cm, Hanging: 1.27 cm, Don't swap indents on facing pages

Deleted: O'Brien, T. A., Risser, M. D., Loring, B., Elbashandy, A. A., Krishnan, H., Johnson, J., ... & Collins, W. D. (2020). Detection of atmospheric rivers with inline uncertainty quantification: TECA-BARD v1. 0.1. *Geoscientific Model Development*, 13(12), 6131-6148.

Deleted: (2020).

Deleted: O'Gorman, P. A., & Dwyer, J. G. (2018). Using machine learning to parameterize moist convection: Potential for modeling of climate, climate change, and extreme events. *Journal of Advances in Modeling Earth Systems*, 10(10), 2548-2563.

Deleted: (2018).

Deleted: O'Leary, J., Paulson, J. A., & Mesbah, A. (2022). Stochastic physics-informed neural ordinary differential equations. *Journal of Computational Physics*, 468, 111466.

Deleted: (2022).

Deleted: (2020).

Deleted: (2021).

Deleted: (2020).

Deleted: (2020).

Deleted: Patel, R. G., Trask, N. A., Wood, M. A., & Cyr, E. C. (2021). A physics-informed operator regression framework for extracting data-driven continuum models. *Computer Methods in Applied Mechanics and Engineering*, 373, 113500.

Deleted: (2021).

Deleted: Peairs, L., & Chen, T. Y. (2011). Using reinforcement learning to vary the m in GMRES (m). *Procedia Computer Science*, 4, 2257-2266.

Deleted: (2011).

Deleted: (2020, September).

Formatted: Not Highlight

31087 [Pincus, R., Mlawer, E. J., & Delamere, J. S. \(2019\). Balancing accuracy, efficiency, and flexibility in radiation calculations for dynamical models. *J ADV MODEL EARTH SY*, 11\(10\), 3074-3089. <https://doi.org/10.1029/2019MS001621>. 2019.](#)

31088

31089

31090 [Prabhat, P., Kashinath, K., Mudigonda, M., Kim, S., Kapp-Schwoerer, L., Graubner, A., & Collins, W. ClimateNet: An expert-labelled open dataset and Deep Learning architecture for enabling high-precision analyses of extreme weather. *GEOSCI MODEL DEV*, 14\(1\), 107-124. <https://doi.org/10.5194/gmd-14-107-2021>. 2021.](#)

31091

31092

31093 [Psaros, A. F., Meng, X., Zou, Z., Guo, L., & Karniadakis, G. E. Uncertainty quantification in scientific machine learning: Methods, metrics, and comparisons. *arXiv preprint arXiv:2201.07766*. <https://doi.org/10.48550/arXiv.2201.07766>. 19 January 2022.](#)

31094

31095

31096 [Rasp, S. Coupled online learning as a way to tackle instabilities and biases in neural network parameterizations: general algorithms and Lorenz 96 case study \(v1. 0\). *GEOSCI MODEL DEV*, 13\(5\), 2185-2196. <https://doi.org/10.48550/arXiv.1907.01351>. 24 March 2020.](#)

31097

31098

31099 [Rasp, S., Dueben, P. D., Scher, S., Weyn, J. A., Mouatadid, S., & Thuerey, N. WeatherBench: a benchmark data set for data-driven weather forecasting. *J ADV MODEL EARTH SY*, 12\(11\), e2020MS002203. <https://doi.org/10.1029/2020MS002203>. 2020.](#)

31100

31101

31102 [Rasp, S., Pritchard, M. S., & Gentine, P. 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>. 2018.](#)

31103

31104 [Rasp, S., & Thuerey, N. Data-driven medium-range weather prediction with a resnet pretrained on climate simulations: A new model for weatherbench. *J ADV MODEL EARTH SY*, 13\(2\), e2020MS002405. <https://doi.org/10.1029/2020MS002405>. 2021.](#)

31105

31106

31107 [Ravuri, S., Lenc, K., Willson, M., Kangin, D., Lam, R., Mirowski, P., ... & Mohamed, S. Skilful precipitation nowcasting using deep generative models of radar. *Nature*, 597\(7878\), 672-677. <https://doi.org/10.1038/s41586-021-03854-z>. 2021.](#)

31108

31109

31110 [Rizzuti, G., Siahkoohi, A., & Herrmann, F. J. Learned iterative solvers for the Helmholtz equation. In *81st EAGE Conference and Exhibition 2019* \(Vol. 2019, No. 1, pp. 1-5\). European Association of Geoscientists & Engineers. <https://doi.org/10.3997/2214-4609.201901542>. 2019.](#)

31111

31112

31113 [Ronneberger, O., Fischer, P., & Brox, T. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention* \(pp. 234-241\). Springer, Cham. \[https://doi.org/10.1007/978-3-319-24574-4_28\]\(https://doi.org/10.1007/978-3-319-24574-4_28\). 2015.](#)

31114

31115

31116 [Rosier, S. H., Bull, C., Woo, W. L., & Gudmundsson, G. H. Predicting ocean-induced ice-shelf melt rates using deep learning. *The Cryosphere*, 17\(2\), 499-518. <https://doi.org/10.5194/tc-17-499-2023>. 2023.](#)

31117

31118 [Ross, A., Li, Z., Perezhogin, P., Fernandez-Granda, C., & Zanna, L. Benchmarking of machine learning ocean subgrid parameterizations in an idealized model. *J ADV MODEL EARTH SY*, 15\(1\), e2022MS003258. <https://doi.org/10.1029/2022MS003258>. 2023.](#)

31119

31120

31121 [Rupe, A., Kashinath, K., Kumar, N., & Crutchfield, J. P. \(2023\). Physics-Informed Representation Learning for Emergent Organization in Complex Dynamical Systems. *arXiv preprint arXiv:2304.12586*. <https://doi.org/10.48550/arXiv.2304.12586>. 25 April 2023.](#)

31122

31123

Deleted: Pincus, R., Mlawer, E. J., & Delamere, J. S. (2019). Balancing accuracy, efficiency, and flexibility in radiation calculations for dynamical models. *Journal of Advances in Modeling Earth Systems*, 11(10), 3074-3089.

Deleted: Prabhat, P., Kashinath, K., Mudigonda, M., Kim, S., Kapp-Schwoerer, L., Graubner, A., & Collins, W. (2021). ClimateNet: An expert-labelled open dataset and Deep Learning architecture for enabling high-precision analyses of extreme weather. *Geoscientific Model Development*, 14(1), 107-124.

Deleted: (2021).

Deleted: (2022).

Deleted: 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.

Deleted: Rasp, S. (2020). Coupled online learning as a way to tackle instabilities and biases in neural network parameterizations: general algorithms and Lorenz 96 case study (v1. 0). *Geoscientific Model Development*, 13(5), 2185-2196.

Deleted: (2020).

Deleted: Rasp, S., Dueben, P. D., Scher, S., Weyn, J. A., Mouatadid, S., & Thuerey, N. (2020). WeatherBench: a benchmark data set for data-driven weather forecasting. *Journal of Advances in Modeling Earth Systems*, 12(11), e2020MS002203.

Deleted: . (2020).

Deleted: (2018).

Deleted: Rasp, S., & Thuerey, N. (2021). Data-driven medium-range weather prediction with a resnet pretrained on climate simulations: A new model for weatherbench. *Journal of Advances in Modeling Earth Systems*, 13(2), e2020MS002405....

Deleted: (2021).

Deleted: (2021).

Deleted: . (2019, June).

Deleted: . (2015, October)

Deleted: (2023).

Deleted: (2023). B

3164 Russell S. & Norvig P. *Artificial Intelligence: A Modern Approach* (Fourth Global Edition). Pearson Education. 2021.

3165 Samek, W., Montavon, G., Lapuschkin, S., Anders, C. J., & Müller, K. R. *Explaining deep neural networks and*

3166 *beyond: A review of methods and applications.* P IEEE, 109(3), 247-278.

3167 <https://doi.org/10.1109/JPROC.2021.3060483>. 2021.

3168 Sawada, Y. *Machine learning accelerates parameter optimization and uncertainty assessment of a land surface model.*

3169 *J GEOPHYS RES-ATMOS*, 125(20), e2020JD032688. <https://doi.org/10.1029/2020JD032688>. 2020.

3170 Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., & Monfardini, G. *The graph neural network model.* IEEE T

3171 *NEURAL NETWORKS*, 20(1), 61-80. <https://doi.org/10.1109/TNN.2008.2005605>. 2008.

3172 Scher, S. *Toward data-driven weather and climate forecasting: Approximating a simple general circulation model with*

3173 *deep learning.* GEOPHYS RES LETT, 45(22), 12-616. <https://doi.org/10.1029/2018GL080704>. 2018.

3174 Scher, S., & Messori, G. *Weather and climate forecasting with neural networks: using general circulation models*

3175 *(GCMs) with different complexity as a study ground.* GEOSCI MODEL DEV, 12(7), 2797-2809.

3176 <https://doi.org/10.5194/gmd-12-2797-2019>. 2019.

3177 Shi, X., Chen, Z., Wang, H., Yeung, D. Y., Wong, W. K., & Woo, W. C. *Convolutional LSTM network: A machine*

3178 *learning approach for precipitation nowcasting.* ADV NEUR IN, 28. 2015.

3179 Taylor, J., & Feng, M. *A Deep Learning Model for Forecasting Global Monthly Mean Sea Surface Temperature*

3180 *Anomalies.* arXiv preprint arXiv:2202.09967. <https://doi.org/10.48550/arXiv.2202.09967>. 21 February

3181 2022.

3182 Tibshirani, R., & Friedman, J. H. *The elements of statistical learning [electronic resource]: data mining, inference,*

3183 *and prediction: with 200 full-color illustrations.* Springer. 2001.

3184 Tompson, J., Schlachter, K., Sprechmann, P., & Perlin, K. *Accelerating eulerian fluid simulation with convolutional*

3185 *networks.* In *International Conference on Machine Learning* (pp. 3424-3433). PMLR. 2017.

3186 Toms, B. A., Barnes, E. A., & Ebert-Uphoff, I. *Physically interpretable neural networks for the geosciences:*

3187 *Applications to earth system variability.* J ADV MODEL EARTH SY, 12(9), e2019MS002002.

3188 <https://doi.org/10.1029/2019MS002002>. 2020.

3189 Turing, A. M., *Computing Machinery and Intelligence*, Mind, Volume LIX, Issue 236, Pages 433-460.

3190 <https://doi.org/10.1093/mind/LIX.236.433>. 1950.

3191 Ukkonen, P., & Mäkelä, A. *Evaluation of machine learning classifiers for predicting deep convection.* J ADV MODEL

3192 *EARTH SY*, 11(6), 1784-1802. <https://doi.org/10.1029/2018MS001561>. 2019.

3193 Ukkonen, P., Pincus, R., Hogan, R. J., Pagh Nielsen, K., & Kaas, E. (2020). *Accelerating radiation computations for*

3194 *dynamical models with targeted machine learning and code optimization.* J ADV MODEL EARTH SY,

3195 *12(12), e2020MS002226.* <https://doi.org/10.1029/2020MS002226>. 2020.

3196 United Nations Educational, Scientific and Cultural Organization. *Recommendations on the Ethics of Artificial*

3197 *Intelligence.* 2021.

3198 Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. *Attention is all you*

3199 *need.* ADV NEUR IN, 30. 2017.

- Deleted: (2021).
- Deleted: Samek, W., Montavon, G., Lapuschkin, S., Anders, C. J., & Müller, K. R. (2021). Explaining deep neural networks and beyond: A review of methods and applications. Proceedings of the IEEE, 109(3), 247-278.
- Deleted: (2021).
- Deleted: (2020).
- Deleted: Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., & Monfardini, G. (2008). The graph neural network model. IEEE transactions on neural networks, 20(1), 61-80.
- Deleted: (2008).
- Deleted: Scher, S. (2018). Toward data-driven weather and climate forecasting: Approximating a simple general circulation model with deep learning. *Geophysical Research Letters*, 45(22), 12-616.
- Deleted: (2018).
- Deleted: Scher, S., & Messori, G. (2019). Weather and climate forecasting with neural networks: using general circulation models (GCMs) with different complexity as a study ground. *Geoscientific Model Development*, 12(7), 2797-2809.
- Deleted: (2019).
- Deleted: Shi, X., Chen, Z., Wang, H., Yeung, D. Y., Wong, W. K., & Woo, W. C. (2015). Convolutional LSTM network: A machine learning approach for precipitation nowcasting. *Advances in neural information processing systems*, 28.
- Deleted: (2015).
- Deleted: Sonderby, C. K., Espelholt, L., Heck, J., Dehghani, M., Oliver, A., Salimans, T., ... & Kalchbrenner, N. (2020). Metnet: A neural weather model for precipitation forecasting. *arXiv preprint arXiv:2003.12140*.
- Deleted: (2022).
- Deleted: (2001).
- Deleted: (2017, July).
- Deleted: Toms, B. A., Barnes, E. A., & Ebert-Uphoff, I. (2020). Physically interpretable neural networks for the geosciences: Applications to earth system variability. (... [29])
- Deleted: (2020). P
- Deleted: , October 1950
- Deleted: Ukkonen, P., & Mäkelä, A. (2019). Evaluati (... [30])
- Deleted: (2019).
- Deleted: Ukkonen, P., Pincus, R., Hogan, R. J., Pagh (... [31])
- Deleted: (2021).
- Deleted: Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30, 6001-6010. (... [32])
- Deleted: (2017).

3265 [Virnodkar, S. S., Pachghare, V. K., Patil, V. C., & Jha, S. K. Remote sensing and machine learning for crop water](#)
3266 [stress determination in various crops: a critical review. PRECIS AGRIC, 21\(5\), 1121-1155.](#)
3267 <https://doi.org/10.1007/s11119-020-09711-9>. 2020.

3268 Vlachas, P. R., Byeon, W., Wan, Z. Y., Sapsis, T. P., & Koumoutsakos, P. Data-driven forecasting of high-dimensional
3269 chaotic systems with long short-term memory networks. *Proceedings of the Royal Society A: Mathematical,*
3270 *Physical and Engineering Sciences*, 474(2213), 20170844. <https://doi.org/10.1098/rspa.2017.0844>. 2018.

3271 [Walters, D., Boutle, I., Brooks, M., Melvin, T., Stratton, R., Vosper, S., ... & Xavier, P. The Met Office unified model](#)
3272 [global atmosphere 6.0/6.1 and JULES global land 6.0/6.1 configurations. GEOSCI MODEL DEV, 10\(4\),](#)
3273 [1487-1520. https://doi.org/10.5194/gmd-10-1487-2017](https://doi.org/10.5194/gmd-10-1487-2017). 2017.

3274 Wang, S., Wang, H., & Perdikaris, P. Learning the solution operator of parametric partial differential equations with
3275 physics-informed DeepONets. *Science advances*, 7(40), eabi8605. <https://doi.org/10.1126/sciadv.abi8605>.
3276 2021.

3277 [Wang, X., Han, Y., Xue, W., Yang, G., & Zhang, G. J. Stable climate simulations using a realistic general circulation](#)
3278 [model with neural network parameterizations for atmospheric moist physics and radiation processes.](#)
3279 [GEOSCI MODEL DEV, 15\(9\), 3923-3940. https://doi.org/10.5194/gmd-15-3923-2022. 2022.](#)

3280 Wang, S., Sankaran, S., & Perdikaris, P. (2022b). Respecting causality is all you need for training physics-informed
3281 neural networks. *arXiv preprint arXiv:2203.07404*. <https://doi.org/10.48550/arXiv.2203.07404>. 14 March
3282 2022.

3283 [Watson, P. A. Machine learning applications for weather and climate need greater focus on extremes. ENVIRON RES](#)
3284 [LETT, 17\(11\), 111004. https://doi.org/10.1088/1748-9326/ac9d4e. 2022.](#)

3285 [Watt-Meyer, O., Brenowitz, N. D., Clark, S. K., Henn, B., Kwa, A., McGibbon, J., ... & Bretherton, C. S. Correcting](#)
3286 [weather and climate models by machine learning nudged historical simulations. GEOPHYS RES LETT,](#)
3287 [48\(15\), e2021GL092555. https://doi.org/10.1029/2021GL092555](https://doi.org/10.1029/2021GL092555). 2021.

3288 [Watson-Parris, D. Machine learning for weather and climate are worlds apart. Philosophical Transactions of the Royal](#)
3289 [Society A, 379\(2194\), 20200098. https://doi.org/10.1098/rsta.2020.0098. 2021.](#)

3290 [Watson-Parris, D., Rao, Y., Oliv  , D., Seland,  ., Nowack, P., Camps-Valls, G., ... & Roesch, C. ClimateBench v1.](#)
3291 [0: A Benchmark for Data-Driven Climate Projections. J ADV MODEL EARTH SY, 14\(10\),](#)
3292 [e2021MS002954. https://doi.org/10.1029/2021MS002954](https://doi.org/10.1029/2021MS002954). 2022.

3293 Werbos, P. Beyond regression: new tools for prediction and analysis in the behavioral sciences. *Ph. D. dissertation,*
3294 *Harvard University. 1974.*

3295 Werbos, P. J. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10), 1550-
3296 1560. <https://doi.org/10.1109/5.58337>. 1990.

3297 [Weyn, J. A., Durran, D. R., & Caruana, R. \(2019\). Can machines learn to predict weather? Using deep learning to](#)
3298 [predict gridded 500-hPa geopotential height from historical weather data. J ADV MODEL EARTH SY,](#)
3299 [11\(8\), 2680-2693. https://doi.org/10.1029/2019MS001705](https://doi.org/10.1029/2019MS001705). 2019.

Deleted: (2020).

Deleted: (2018).

Deleted: Walters, D., Boutle, I., Brooks, M., Melvin, T., Stratton, R., Vosper, S., ... & Xavier, P. (2017). The Met Office unified model global atmosphere 6.0/6.1 and JULES global land 6.0/6.1 configurations. *Geoscientific Model Development*, 10(4), 1487-1520.

Deleted: (2017).

Deleted: (2021).

Deleted: Wang, X., Han, Y., Xue, W., Yang, G., & Zhang, G. J. (2022a). Stable climate simulations using a realistic general circulation model with neural network parameterizations for atmospheric moist physics and radiation processes. *Geoscientific Model Development*, 15(9), 3923-3940.

Deleted:

Deleted: (2022a).

Deleted: Watson, P. A. (2022). Machine learning applications for weather and climate need greater focus on extremes. *Environmental Research Letters*, 17(11), 111004.

Deleted: (2022).

Formatted: Indent: Left: 0 cm, Hanging: 2.54 cm, Don't swap indents on facing pages

Deleted: Watt-Meyer, O., Brenowitz, N. D., Clark, S. K., Henn, B., Kwa, A., McGibbon, J., ... & Bretherton, C. S. (2021). Correcting weather and climate models by machine learning nudged historical simulations. *Geophysical Research Letters*, 48(15), e2021GL092555.

Deleted: (2021).

Formatted: Indent: Left: 0 cm, Hanging: 1.27 cm, Don't swap indents on facing pages

Deleted: (2021).

Deleted: Watson-Parris, D., Rao, Y., Oliv  , D., Seland,  ., Nowack, P., Camps-Valls, G., ... & Roesch, C. (2022). ClimateBench v1. 0: A Benchmark for Data-Driven Climate Projections. *Journal of Advances in Modeling Earth Systems*, 14(10), e2021MS002954.

Deleted: . (2022)

Deleted: (1974).

Deleted: (1990).

Deleted: Weyn, J. A., Durran, D. R., & Caruana, R. (2019). Can machines learn to predict weather? Using deep learning to predict gridded 500-hPa geopotential height from historical weather data. *Journal of Advances in Modeling Earth Systems*, 11(8), 2680-2693.

3342 [Weyn, J. A., Durran, D. R., & Caruana, R. Improving data-driven global weather prediction using deep convolutional](#)
3343 [neural networks on a cubed sphere. J ADV MODEL EARTH SY, 12\(9\), e2020MS002109.](#)
3344 <https://doi.org/10.1029/2020MS002109>. 2020.

3345 [Weyn, J. A., Durran, D. R., Caruana, R., & Cresswell-Clay, N. Sub-seasonal forecasting with a large ensemble of](#)
3346 [deep-learning weather prediction models. J ADV MODEL EARTH SY, 13\(7\).](#)
3347 <https://doi.org/10.1029/2021MS002502>. 2021.

3348 Wikner, A., Pathak, J., Hunt, B., Girvan, M., Arcomano, T., Szunyogh, I., ... & Ott, E. Combining machine learning
3349 with knowledge-based modeling for scalable forecasting and subgrid-scale closure of large, complex,
3350 spatiotemporal systems. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 30(5), 053111.
3351 <https://doi.org/10.1063/5.0005541>. 2020.

3352 [Wu, K., & Xiu, D. Data-driven deep learning of partial differential equations in modal space. J COMPUT PHYS, 408,](#)
3353 [109307. https://doi.org/10.1016/j.jcp.2020.109307. 2020.](#)

3354 Yamada, K., Katagiri, T., Takizawa, H., Minami, K., Yokokawa, M., Nagai, T., & Ogino, M. Preconditioner auto-
3355 tuning using deep learning for sparse iterative algorithms. In *2018 Sixth International Symposium on*
3356 *Computing and Networking Workshops (CANDARW)* (pp. 257-262). IEEE.
3357 <https://doi.org/10.1016/j.jcp.2020.109307>. 2018.

3358 [Yang, C., Yang, X., & Xiao, X. Data-driven projection method in fluid simulation. COMPUT ANIMAT VIRT W,](#)
3359 [27\(3-4\), 415-424. https://doi.org/10.1002/cav.1695. 2016.](#)

3360 [Yeo, K., Grullon, D. E., Sun, F. K., Boning, D. S., & Kalagnanam, J. R. Variational inference formulation for a model-](#)
3361 [free simulation of a dynamical system with unknown parameters by a recurrent neural network. SIAM J SCI](#)
3362 [COMPUT, 43\(2\), A1305-A1335. https://doi.org/10.1137/20M1323151. 2021.](#)

3363 [Yuan, Z., Zhou, X., & Yang, T. Hetero-convlstm: A deep learning approach to traffic accident prediction on](#)
3364 [heterogeneous spatio-temporal data. In Proceedings of the 24th ACM SIGKDD International Conference on](#)
3365 [Knowledge Discovery & Data Mining \(pp. 984-992\). https://doi.org/10.1145/3219819.3219922. 2018.](#)

3366 [Yuval, J., & O'Gorman, P. A. \(2020\). Stable machine-learning parameterization of subgrid processes for climate](#)
3367 [modeling at a range of resolutions. NAT COMMUN, 11\(1\), 1-10. https://doi.org/10.1038/s41467-020-](#)
3368 [17142-3. 2020.](#)

3369 [Yuval, J., O'Gorman, P. A., & Hill, C. N. Use of neural networks for stable, accurate and physically consistent](#)
3370 [parameterization of subgrid atmospheric processes with good performance at reduced precision. GEOPHYS](#)
3371 [RES LETT, 48\(6\), e2020GL091363. https://doi.org/10.1029/2020GL091363. 2021.](#)

3372 [Zagoruyko, S., & Komodakis, N. Wide residual networks. arXiv preprint arXiv:1605.07146.](#)
3373 <https://doi.org/10.48550/arXiv.1605.07146>. 23 May 2016.

3374 [Zanna, L., & Bolton, T. Data-driven equation discovery of ocean mesoscale closures. GEOPHYS RES LETT, 47\(17\),](#)
3375 [e2020GL088376. https://doi.org/10.1029/2020GL088376. 2020.](#)

3376 [Zhang, N., Zhou, X., Kang, M., Hu, B. G., Heuvelink, E., & Marcelis, L. F. Machine learning versus crop growth](#)
3377 [models: an ally, not a rival. AOB PLANTS, 15\(2\), plac061. https://doi.org/10.1093/aobpla/plac061. 2023.](#)

Deleted: 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.

Deleted: (2020).

Deleted: Weyn, J. A., Durran, D. R., Caruana, R., & Cresswell-Clay, N. (2021). Sub-seasonal forecasting with a large ensemble of deep-learning weather prediction models. *Journal of Advances in Modeling Earth Systems*, 13(7)

Deleted: (2021).

Deleted: (2020).

Deleted: Wu, K., & Xiu, D. (2020). Data-driven deep learning of partial differential equations in modal space. *Journal of Computational Physics*, 408, 109307.

Deleted: (2020).

Deleted: (2018, November).

Deleted: Yang, C., Yang, X., & Xiao, X. (2016). Data-driven projection method in fluid simulation. *Computer Animation and Virtual Worlds*, 27(3-4), 415-424.

Deleted: (2016).

Deleted: Yeo, K., Grullon, D. E., Sun, F. K., Boning, D. S., & Kalagnanam, J. R. (2021). Variational inference formulation for a model-free simulation of a dynamical system with unknown parameters by a recurrent neural network. *SIAM Journal on Scientific Computing*, 43(2), A1305-A1335.

Deleted: (2021).

Deleted: (2018, July).

Deleted: Yuval, J., & O'Gorman, P. A. (2020). Stable machine-learning parameterization of subgrid processes for climate modeling at a range of resolutions. *Nature communications*, 11(1), 1-10.

Deleted: Yuval, J., O'Gorman, P. A., & Hill, C. N. (2021). Use of neural networks for stable, accurate and physically consistent parameterization of subgrid atmospheric processes with good performance at reduced precision. *Geophysical Research Letters*, 48(6), e2020GL091363.

Deleted: (2021).

Deleted: (2016).

Deleted: Zanna, L., & Bolton, T. (2020). Data-driven equation discovery of ocean mesoscale closures. *Geophysical Research Letters*, 47(17), e2020GL088376.

Deleted: (2020).

Deleted: (2023).

Deleted:

3426 [Zhao, W. L., Gentine, P., Reichstein, M., Zhang, Y., Zhou, S., Wen, Y., ... & Qiu, G. Y. Physics-constrained machine](#)
 3427 [learning of evapotranspiration. GEOPHYS RES LETT, 46\(24\), 14496-14507.](#)
 3428 <https://doi.org/10.1029/2019GL085291>. 2019.

3429 [Zhong, X., Ma, Z., Yao, Y., Xu, L., Wu, Y., & Wang, Z. WRF-ML v1.0: a bridge between WRF v4.3 and machine](#)
 3430 [learning parameterizations and its application to atmospheric radiative transfer. GEOSCI MODEL DEV,](#)
 3431 [16\(1\), 199-209. https://doi.org/10.5194/gmd-16-199-2023. 2023.](#)

3432 [Zhou, L., Lin, S. J., Chen, J. H., Harris, L. M., Chen, X., & Rees, S. L. Toward convective-scale prediction within the](#)
 3433 [next generation global prediction system. B AM METEOROL SOC, 100\(7\), 1225-1243.](#)
 3434 <https://doi.org/10.1175/BAMS-D-17-0246.1>. 2019.

3435

Deleted: (2019).

Deleted: Zhong, X., Ma, Z., Yao, Y., Xu, L., Wu, Y., & Wang, Z. (2023). WRF-ML v1.0: a bridge between WRF v4.3 and machine learning parameterizations and its application to atmospheric radiative transfer. *Geoscientific Model Development*, 16(1), 199-209.

Deleted: (2023).

Deleted: Zhou, L., Lin, S. J., Chen, J. H., Harris, L. M., Chen, X., & Rees, S. L. (2019). Toward convective-scale prediction within the next generation global prediction system. *Bulletin of the American Meteorological Society*, 100(7), 1225-1243.

Deleted: (2019).

Page 6: [1] Deleted Catherine de Burgh-Day (she/her) 18/07/2023 19:28:00

Page 10: [2] Deleted Cat Cat 21/07/2023 12:54:00

Page 10: [3] Deleted Catherine de Burgh-Day (she/her) 18/07/2023 18:58:00

Page 10: [4] Deleted Cat Cat 21/07/2023 12:55:00

Page 40: [5] Deleted Catherine de Burgh-Day (she/her) 13/07/2023 23:38:00

Page 40: [6] Deleted Catherine de Burgh-Day (she/her) 14/07/2023 01:39:00

Page 42: [7] Deleted Tennessee Leeuwenburg (he/him) 13/07/2023 17:27:00

Page 42: [7] Deleted Tennessee Leeuwenburg (he/him) 13/07/2023 17:27:00

Page 42: [8] Deleted Tennessee Leeuwenburg (he/him) 28/06/2023 00:46:00

Page 42: [9] Deleted Catherine de Burgh-Day (she/her) 21/07/2023 10:25:00

Page 42: [9] Deleted Catherine de Burgh-Day (she/her) 21/07/2023 10:25:00

Page 42: [10] Deleted Catherine de Burgh-Day (she/her) 21/07/2023 10:27:00

Page 42: [10] Deleted Catherine de Burgh-Day (she/her) 21/07/2023 10:27:00

Page 42: [11] Deleted Catherine de Burgh-Day (she/her) 21/07/2023 10:28:00

Page 42: [11] Deleted Catherine de Burgh-Day (she/her) 21/07/2023 10:28:00

Page 42: [11] Deleted Catherine de Burgh-Day (she/her) 21/07/2023 10:28:00

Page 42: [11] Deleted Catherine de Burgh-Day (she/her) 21/07/2023 10:28:00

Page 42: [11] Deleted Catherine de Burgh-Day (she/her) 21/07/2023 10:28:00

Page 42: [11] Deleted Catherine de Burgh-Day (she/her) 21/07/2023 10:28:00

Page 42: [12] Deleted Catherine de Burgh-Day (she/her) 21/07/2023 10:40:00

Page 42: [12] Deleted Catherine de Burgh-Day (she/her) 21/07/2023 10:40:00

Page 42: [12] Deleted Catherine de Burgh-Day (she/her) 21/07/2023 10:40:00

Page 42: [12] Deleted Catherine de Burgh-Day (she/her) 21/07/2023 10:40:00

Page 43: [13] Deleted Catherine de Burgh-Day (she/her) 21/07/2023 10:42:00

Page 43: [14] Deleted Catherine de Burgh-Day (she/her) 21/07/2023 10:58:00

Page 43: [15] Deleted Catherine de Burgh-Day (she/her) 21/07/2023 10:59:00

Page 43: [16] Deleted Catherine de Burgh-Day (she/her) 22/06/2023 06:48:00

Page 43: [17] Deleted Catherine de Burgh-Day (she/her) 22/06/2023 06:23:00

Page 47: [18] Deleted Cat Cat 21/07/2023 13:00:00

Page 61: [19] Deleted Tennessee Leeuwenburg (he/him) 20/07/2023 11:33:00

Page 61: [20] Formatted Cat Cat 21/07/2023 13:55:00

Indent: Left: 0 cm, Hanging: 1.27 cm, Don't swap indents on facing pages

Page 61: [21] Formatted Catherine de Burgh-Day (she/her) 20/07/2023 16:19:00

Indent: Left: 0 cm, Hanging: 1.27 cm, Don't swap indents on facing pages

Page 61: [22] Deleted Tennessee Leeuwenburg (he/him) 20/07/2023 11:34:00

Page 61: [23] Deleted Tennessee Leeuwenburg (he/him) 20/07/2023 11:34:00

Page 61: [24] Deleted Tennessee Leeuwenburg (he/him) 20/07/2023 11:35:00

Page 61: [25] Deleted Tennessee Leeuwenburg (he/him) 20/07/2023 11:36:00

Page 61: [26] Deleted Tennessee Leeuwenburg (he/him) 20/07/2023 11:36:00

Page 64: [27] Deleted Tennessee Leeuwenburg (he/him) 20/07/2023 11:57:00

Page 64: [28] Deleted Tennessee Leeuwenburg (he/him) 20/07/2023 11:58:00

Page 70: [29] Deleted **Tennessee Leeuwenburg (he/him)** **20/07/2023 13:01:00**

Page 70: [30] Deleted **Tennessee Leeuwenburg (he/him)** **20/07/2023 13:01:00**

Page 70: [31] Deleted **Tennessee Leeuwenburg (he/him)** **20/07/2023 13:02:00**

Page 70: [32] Deleted **Tennessee Leeuwenburg (he/him)** **20/07/2023 13:02:00**