

# 1 **Deep Learning, Explained:** 2 **Fundamentals, Explainability, and Bridgeability to Process-based Modelling**

3 Saman Razavi

4 Associate Professor, Global Institute for Water Security, School of Environment and Sustainability,  
5 University of Saskatchewan, Department of Civil, Geological and Environmental Engineering, Canada  
6 [saman.razavi@usask.ca](mailto:saman.razavi@usask.ca)

## 7 **Abstract**

8 Recent breakthroughs in artificial intelligence (AI), and particularly in deep learning (DL), have created  
9 tremendous excitement and opportunities in the earth and environmental sciences communities. To  
10 leverage these new 'data-driven' technologies, however, one needs to understand the fundamental  
11 concepts that give rise to DL and how they differ from 'process-based', mechanistic modelling. This paper  
12 revisits those fundamentals and addresses 10 questions often posed by earth and environmental  
13 scientists with the aid of a real-world modelling experiment. The overarching objective is to contribute to  
14 a future of AI-assisted earth and environmental sciences where DL models can (1) embrace the typically  
15 ignored knowledge base available, (2) function credibly in 'true' out-of-sample prediction, and (3) handle  
16 non-stationarity in earth and environmental systems. Comparing and contrasting earth and  
17 environmental problems with prominent AI applications, such as playing chess and trading in the stock  
18 market, provides critical insights for better directing future research in this field.

## 19 **Plain Language Summary**

20 Deep learning (DL) is an artificial intelligence (AI) technique that has already served the vast majority, if  
21 not all, of everyday society in tasks such as image recognition and language processing through  
22 smartphones. The recent unprecedented performance of DL in those tasks has accelerated applications  
23 in non-native areas such as earth and environmental sciences where knowledge-based modelling has  
24 dominated to date. A major challenge, however, is DL and knowledge-based modelling are rooted in  
25 different worldviews towards problem solving. This paper explains the 'whats' and 'whys' of DL from first  
26 principles, with an eye on applications since inception in environmental problems. An experiment is run  
27 to illustrate the fundamental differences between the two worldviews, and to shed light on some critical,  
28 but often ignored, issues DL may face in practice, largely arising from the fact that earth and  
29 environmental systems are complex with behaviors changing in ways that are physically explainable but  
30 not seen in the period of record due to uncertain factors such as climate change. Such issues must be  
31 addressed at the heart of the endeavor to develop DL techniques that embrace the knowledge base  
32 available, in anticipation of breakthroughs in an age of big data and computational power.

## 33 **Keywords**

34 Artificial intelligence, machine learning, deep learning, artificial neural networks, process-based  
35 modelling, earth systems, hydrology

36

37 **Key Points**

- 38 • DL is rooted in *connectionism*, *hyper-flexibility*, and *vigorous optimization*, which are alien to  
 39 conventional knowledge-based modelling.
- 40 • A knowledge base is essential to enable credible predictions of *complex*, *open*, *partially observable*,  
 41 and *non-stationary systems*.
- 42 • Bridging DL and earth and environmental sciences is still embryonic but has great potential in an age  
 43 of big data and computational power.

44 **Table of Contents**

45 1. The rise of deep learning ..... 3

46 2. Back to fundamentals ..... 4

47 2.1. Why AI and DL? ..... 4

48 2.2. Evolution of DL and major milestones ..... 5

49 2.3. Latest developments and rebranding the field ..... 7

50 3. Geometrical Interpretation of DL ..... 9

51 3.1. A perceptron ..... 9

52 3.2. ANNs with one hidden layer ..... 10

53 3.3. Why more than one hidden layer? ..... 11

54 4. Relevance of Occam’s razor and equifinality? ..... 13

55 4.1. Issues with the complexity of ANNs ..... 13

56 4.2. Leashing the hyper-flexibility of ANNs ..... 14

57 5. Fundamental differences from other ML methods ..... 16

58 5.1. Local versus distributed representations ..... 16

59 5.2. Implications for users ..... 16

60 6. How to introduce order, time-dependency, and memory ..... 17

61 6.1. Tapped delay lines ..... 17

62 6.2. Recurrent connections ..... 18

63 6.3. Training considerations when the order of data matters ..... 19

64 7. ML versus process-based modelling – An experiment ..... 20

65 7.1. Data and models ..... 20

66 7.2. Model performance in calibration ..... 22

67 7.3. What about *a priori* information encoded in models? ..... 23

68 7.4. Model validation: Standard versus true out-of-sample prediction ..... 23

69 7.5. Injecting some physics into ML ..... 25

70 7.6. So, what model should we trust: the ML or physically based model? ..... 26

71 8. Discussion ..... 26

72 8.1. What is the typically ignored value of domain knowledge in DL? ..... 26

73 8.2. Why is DL essentially different from process-based modelling? ..... 28

74 8.3. How can we bridge DL and process-based modelling? ..... 28

75 8.4. What can we learn from prominent DL applications? ..... 30

76 9. Concluding remarks ..... 31

77 Acknowledgements ..... 32

78 References ..... 32

79

80 **1. The rise of deep learning**

81 The last decade has witnessed a tremendous rise in techniques called ‘deep learning’ (DL), under the  
82 umbrella of artificial intelligence (AI) and machine learning (ML), and their unprecedented performance  
83 in areas such as computer vision (Krizhevsky et al., 2017), natural language processing (Young et al., 2018),  
84 and gaming (Silver et al., 2018). These successes have motivated the application of DL across a wide range  
85 of disciplines, including medicine (Hosny et al., 2018), earth sciences (Reichstein et al., 2019), robotics  
86 (Torresen, 2018), engineering (Panchal et al., 2019), and finance (Lee et al., 2019). DL owes its exemplary  
87 success to the boom in computational power and the emergence of big data sources and associated data  
88 storage and sharing technologies.

89 Earth and environmental sciences appear to be positioned to benefit profoundly from DL, as big data  
90 sources on a range of *in situ* and remotely-sensed variables are becoming increasingly available with the  
91 advances in sensing technologies (Reichstein et al., 2019). The storage volume of remote sensing data for  
92 earth observations is already well beyond dozens of petabytes, with transmission rates exceeding  
93 hundreds of terabytes per day. Datasets based on model outputs are rising; for example, the climate  
94 assessment dataset provided by the Coupled Model Intercomparison Project Phase 6 may reach 40  
95 petabytes (Eyring et al., 2016). Reanalysis climatic datasets have also grown; for example, NASA’s Modern-  
96 Era Retrospective Analysis for Research and Applications version 2 (MERRA-2) is ~400 terabytes (Gelaro  
97 et al., 2017). In addition, datasets generated via tens of thousands of citizen science projects are providing  
98 large and rich sources of ground-based data.

99 This potential is shifting the attention of earth and environmental scientists and relevant funding agencies  
100 towards ML, as evidenced, for example, by the shift in research work presented at the American  
101 Geophysical Union (AGU)’s fall meetings, the largest assembly of earth and environmental scientists with  
102 more than 27,000 people in attendance and 25,000 presentations in 2019. The number of ML-related  
103 presentations has risen consistently—from 0.2% of total presentations in 2015 to 4.2% in 2020. In  
104 particular, this shift has been astonishing in the ‘non-linear geophysics’, ‘earth and space science  
105 informatics’, ‘natural hazards’, ‘hydrology’, and ‘seismology’ sub-fields, where 28 (2.1), 18 (5.1), 9 (1.3),  
106 7.5 (1.4), and 6.7% (0.9%) of total presentations, respectively, were related to ML in 2020 (2015).

107 Recent successful applications of DL techniques to earth and environmental sciences include weather  
108 forecasting (Xingjian et al., 2015), rainfall-runoff modelling (Kratzert et al., 2018), rain and snow retrieval  
109 from spaceborne sensors (Tang et al., 2018), downscaling hydroclimatic variables (Ducournau and Fablet,  
110 2016), and surrogate modelling (Razavi et al., 2012a). Unsuccessful applications, perhaps similar to many  
111 other areas, remain largely unreported in the peer-reviewed scientific literature but occasionally appear  
112 in other media (e.g., Wexler, 2017; Kolakowski, 2018).

113 Notably, most DL algorithms, formerly known as artificial neural networks (ANNs), have been around and  
114 widely applied in earth and environmental sciences since the early 1990s. These applications are  
115 documented in reviews by Gardner and Dorling (1998), Maier and Dandy (2000), Krasnopolsky (2007),  
116 Maier et al. (2010), Abrahart et al. (2012), Razavi et al. (2012a), Shen (2018), Bergen et al. (2019), and  
117 Reichstein et al. (2019). Arguably, however, the uptake of DL to facilitate and advance earth and  
118 environmental sciences has not kept pace with data availability and computational power over the past  
119 three decades.

120 But why? The challenges impeding the widespread application of DL to earth and environmental problems  
121 to date may be rooted in the fact that convincingly casting those problems, for which an extensive  
122 knowledge base is usually available, within the DL framework is often not straightforward. Moreover, the  
123 lack of interpretability and explainability of DL has been a major hindrance, as model developers need to  
124 be able to make sense of why a model functions the way it does, and to explain that to model users. These  
125 challenges can be further complicated in the absence of a solid understanding of the fundamentals of DL  
126 and how they differ from theory-driven, mechanistic modelling and prediction.

127 And why this paper? Motivated by the recent breakthroughs by DL in its original areas of application,  
128 namely computer vision and natural language processing, this paper aims to address the persistent  
129 challenges facing DL applications in non-native areas related to earth and environmental sciences. With  
130 this overarching aim, this paper addresses 10 questions regarding the fundamentals of DL and its  
131 explainability and bridgeability to earth and environmental systems modelling:

- 132 (1) What is DL and how did it evolve from ANNs?
- 133 (2) Can we really interpret the internal functioning of DL?
- 134 (3) How can the complexity of DL be justified in light of the principle of parsimony?
- 135 (4) Why is DL considered superior to other types of ML?
- 136 (5) How can DL account for memory and time dependency?
- 137 (6) How do DL and process-based models behave differently in out-of-sample prediction?
- 138 (7) What is the typically ignored value of domain knowledge in DL?
- 139 (8) Why is DL essentially different from process-based modelling?
- 140 (9) What are the existing approaches to bridging DL and mechanistic modelling?
- 141 (10) What can we learn from prominent DL applications such as gaming and the stock market?

142 The structure of this paper is such that it best serves the reader when all sections are followed  
143 sequentially. However, an advanced reader could directly refer to a section designated to address a  
144 question of interest. Sections 2 through 7 address questions 1 through 6 and sub-sections 8.1 through 8.4  
145 address questions 7 through 10, respectively. The contents of this paper are intended to be accessible to  
146 a wide audience from various fields under the umbrella of earth and environmental sciences. However,  
147 the views presented mainly arise from the author's data- and theory-driven research background in  
148 hydrology and water resources. Further, a real-world hydrological modelling problem and multiple  
149 synthetic functions are used to explain complex concepts via simple examples.

## 150 **2. Back to fundamentals**

### 151 **2.1. Why AI and DL?**

152 AI, and in particular DL, is nowadays concerned with developing machines that improve their own  
153 performance in carrying out a given task over time by 'learning' from examples, with minimal human  
154 efforts to instruct the machines how to do so (Jordan and Mitchell, 2015). According to Goodfellow et al.  
155 (2016), however, the early efforts to generate AI were based on a knowledge base paradigm to instruct  
156 machines with a formal set of step-by-step mathematical and if-then rules. Those efforts focused on  
157 carrying out tasks that were intellectually difficult for humans but straightforward for computers.  
158 Goodfellow et al. (2016) argue such efforts led to no major successes, and the AI of today is about enabling  
159 machines to perform tasks that humans perform intuitively and rather easily but have difficulty formally

160 describing how they do so. Examples of such tasks include recognizing faces in a photo or comprehending  
161 spoken words.

162 Not only did state-of-the-art AI divorce from the knowledge base, but it also completely separated from  
163 classic data-driven modelling rooted in statistics such as regression. This separation was a response to the  
164 need for models that are not constrained by the many assumptions typical statistical models hold. For  
165 example, traditional statistical modelling requires a formalization of relationships between variables and  
166 assumptions about functional shapes, distributions of variables, and their inter-dependencies, which  
167 enables hypothesis testing and the generation of confidence bounds (see Dangeti, 2017, p. 10-11).  
168 Conversely, in the DL context the underlying relationships in data may have any complex form, which is  
169 typically unknown *a priori*, and the data used may have any size and distributional properties.

170 Because of these characteristics, DL is deemed suitable to pursue the longstanding ambition to build  
171 machines that work with minimal or no human supervision and imposed assumptions. As a result, DL  
172 provides hyper-flexible tools that can adapt to a wide range of data and applications.

## 173 **2.2. Evolution of DL and major milestones**

174 It was 1957 when Frank Rosenblatt invented the first algorithm, termed ‘perceptron’ (Rosenblatt, 1957),  
175 which today is considered the smallest computational unit of DL. A perceptron, alternatively termed a  
176 ‘neuron’ because of its resemblance to the basic working unit of the brain, is shown in **Figure 1a** and  
177 formulated as:

$$178 \quad y = f\left(\sum_{i=1}^D w_i x_i + b\right) \quad (\text{Eq. 1})$$

179 where  $D$  is the dimension of input space,  $\mathbf{x}$  is the input vector,  $\mathbf{w}$  is a set of weights corresponding to the  
180 input vector,  $b$  is bias, and  $f$  is an ‘activation’ function. A perceptron has  $D+1$  tunable parameters (i.e.,  $D$   
181 weights and one bias) and is basically nothing but a multiple linear regression augmented by an output  
182 function ( $f$ ), which is typically non-linear. The form of the activation function was originally a step function,  
183 but now a range of monotonic functional forms, most commonly ‘sigmoidal’, are used.

184 The invention of perceptrons created significant excitement in the AI community and beyond. But, it soon  
185 became clear that a perceptron would not be able to map input spaces that are not linearly separable,  
186 such as the XOR problem (Minsky and Papert, 1969), rendering perceptrons of limited use in real-world  
187 applications. The reason for this inability is that the core of the perceptron is a linear regression.

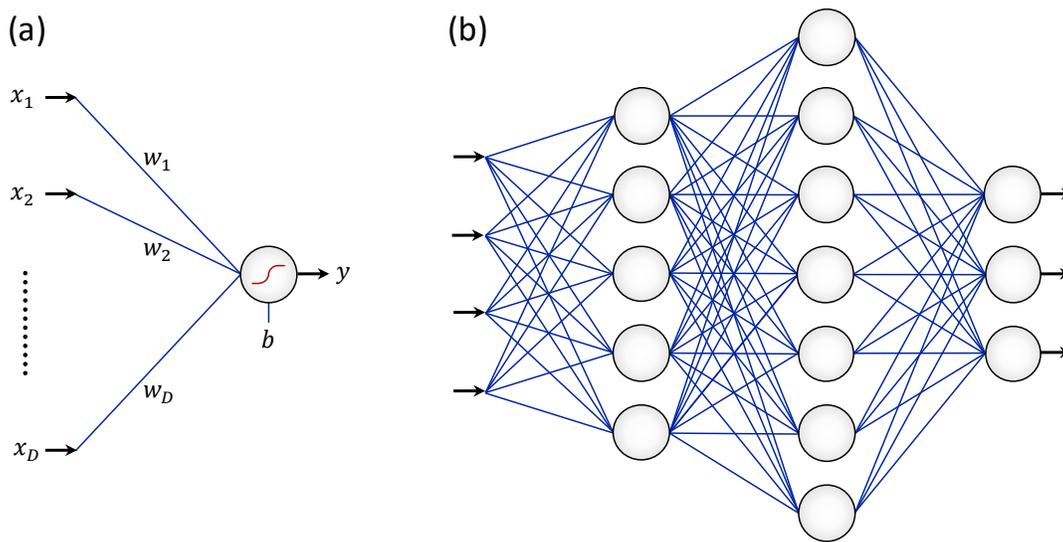
188 Efforts to overcome this barrier could have followed two different avenues. Perhaps the most intuitive  
189 avenue was to employ non-linear regression, by allowing the terms inside the parentheses in Eq. 1 to be  
190 of other algebraic forms. However, this was not a viable option in part because the user then would need  
191 to specify the form of non-linearity *a priori*, which was not compatible with the principles of AI.

192 The second avenue that led to today’s DL was to combine perceptrons both in parallel and in series to  
193 create so-called ‘multi-layer perceptrons’ (MLPs), as shown in **Figure 1b**, with the hope this more complex  
194 system could overcome the barrier. An MLP would then have many more tunable parameters than the  
195 perceptron. The first layer, also called the first ‘hidden’ layer, would have  $n_1 \cdot D$  weights and  $n_1$  biases,  
196 where  $n_1$  is the number of neurons in this layer. Similarly, the second hidden layer would have  $n_2 \cdot n_1$   
197 weights and  $n_2$  biases, and the last layer, called the ‘output’ layer would have  $n_d \cdot n_{d-1}$  weights and  $n_d$  biases,

198 where  $n_2$  and  $n_d$  are the numbers of neurons in the second and last layers, respectively, and  $d$  is the total  
199 number of layers.

200 The total number of layers in an MLP and the number of neurons in each layer are ‘hyper-parameters’, to  
201 be specified by users. Also important is the choice of the activation function in each layer. Note that a  
202 linear activation function is typically only suitable for the last layer and, in general, any stack of linear  
203 layers is effectively equivalent to a single linear layer. MLPs have also historically been called ‘artificial  
204 neural networks’ (ANNs), or simply ‘neural networks’ (NNs), because of their perceived resemblance to  
205 biological neural networks.

206



207

208 **Figure 1.** (a) A perceptron and (b) a multi-layer perceptron with four inputs, two hidden layers, and  
209 three outputs.

210

211 MLPs on their own did not go far and the field stagnated for many years because of the absence of an  
212 algorithm that could automatically derive from data the network weights and biases—a process referred  
213 to as ‘training’ in the AI community. It took until the mid-1980s when the first ‘back-propagation’ (BP)  
214 algorithm was invented to enable the training of MLPs with any network structure (Rumelhart et al.,  
215 1986). This invention marked the beginning of the ‘second wave’ of popularity of ANNs. BP is essentially  
216 an optimization algorithm, based on non-linear programming, that minimizes a loss function representing  
217 the goodness-of-fit of predictions to observations, such as the ‘sum of squared errors’, as follows:

$$218 \quad F = \sum_{k=1}^M \sum_{j=1}^N (T_j^k - y_j^k)^2 \quad (\text{Eq. 2})$$

219 where  $y_j^k$  is the output of neuron  $j$  in the output layer when the network is forced with input data entry  $k$   
220 and  $T_j^k$  is the respective desired target. Also,  $M$  is the size of training data, and  $N$  is the number of neurons  
221 in the output layer.

222 Different variations of BP rooted in first- or second-order optimization, or a combination thereof, now  
223 exist; see e.g., the Levenberg-Marquardt algorithm as implemented by Hagan and Menhaj (1994). These  
224 algorithms are fundamentally the same as optimization algorithms used nowadays for calibration of  
225 process-based models. The only difference is that, in the case of ANNs, and unlike most process-based  
226 models, the partial derivatives of the loss function with respect to weights and biases are *analytically*  
227 available and obtained through the ‘chain rule of differentiation’. More recently, derivative-free and  
228 metaheuristic optimization algorithms have shown promise in ANN training (e.g., Dengiz et al., 2009;  
229 Rakitianskaia and Engelbrecht, 2009; Razavi and Tolson, 2011), but have yet to become mainstream.

230 In the late 1980s, after the invention of BP, MLPs were proven to be ‘universal approximators’ (Hornik et  
231 al., 1989). This proof indicated MLPs with only one single-hidden layer that possesses a sigmoidal  
232 activation function, and a linear output layer, would be able to approximate any function with any desired  
233 level of accuracy provided the number of hidden neurons is sufficient. Since then, the ‘universal function  
234 approximation theorem’ has been the fundamental driver of interest in MLPs across a variety of disciplines  
235 and applications.

236 ANNs started receiving much attention in earth and environmental sciences in the early 1990s. The  
237 pioneering applications of ANNs include: Benediktsson et al. (1990), Badran et al. (1991), Stogryn et al.  
238 (1994), Bankert (1994), and Cabrera-Mercader and Staelin (1995) in the context of remote sensing of the  
239 environment; McCann (1992), Boznar et al. (1993), and Navone and Ceccatto (1994) in the context of  
240 atmospheric forecasting; and Kang et al. (1993), Hsu et al. (1995), and Minns and Hall (1996) in the context  
241 of hydrology modelling. Perhaps the most prominent and widely used application of ANNs in these fields  
242 has been related to the development of PERSIANN, or ‘Precipitation Estimation from Remotely Sensed  
243 Information using Artificial Neural Networks’ (Hsu et al., 1997; Sorooshian et al., 2000; Ashouri et al.,  
244 2015), which has been maintained and updated for two decades (accessible at  
245 <https://chrsdata.eng.uci.edu/>).

246 Despite all of these advances, investments in ANNs and therefore the popularity of ANNs saw a decline in  
247 the AI community beginning in the mid-1990s, perhaps triggered by failures to fulfill overly ambitious or  
248 unrealistic promises by prominent AI scientists (Goodfellow et al., 2016), as historically observed in ‘AI  
249 winters’ (Hendler, 2008). ANNs in earth and environmental sciences, however, remained fairly popular  
250 arguably until the mid-2000s. The focus of researchers in these fields was to find novel applications of  
251 ANNs across different earth and environmental problems.

252 It took until early 2010s before the third wave of popularity and interest in ANNs hit, when the field was  
253 revived and renamed ‘deep learning’. ‘Depth’ is a recently popularized term and loosely refers to the  
254 number of hidden layers in ANNs. A related term is ‘width’, which loosely refers to the number of neurons  
255 in hidden layers. Now, a DL model (or deep ANN) simply refers to an MLP with two or more hidden layers.  
256 All of the recent excitement around ANNs is despite the fact that the structure, formulation, and other  
257 properties of MLPs have remained unchanged since their inception, except for some minor modifications.  
258 So, one might ask: is DL merely a repackaging and rebranding of what existed before? The next section  
259 attempts to answer this question while reviewing the recent milestones.

### 260 **2.3. Latest developments and rebranding the field**

261 To better understand the recent developments in ANNs, one first needs to know the history around the  
262 ‘depth’ concept. MLPs, since their inception, have been used with various numbers of hidden layers, that

263 is with various depths. Most applications, however, remained limited to networks with only one hidden  
264 layer until very recently. For example, [Razavi et al. \(2012a\)](#) report that more than 90% of ANNs used for  
265 surrogate modelling in water resources literature have only one hidden layer. There was (and perhaps still  
266 is) no consensus about a proper network depth, because identifying the optimal network configuration  
267 for a given problem and dataset is challenging.

268 Historically, some researchers favored ANNs with more than one hidden layer, arguing that they require  
269 fewer hidden neurons to approximate the same function (see e.g., [Tamura and Tateishi, 1997](#)). On the  
270 other hand, others asserted that single-hidden-layer ANNs are superior to those with more than one  
271 hidden layer with the same level of complexity (see e.g., [de Villiers and Barnard, 1993](#)). A discussion on  
272 this matter is available in [Razavi et al. \(2012a\)](#).

273 Three general reasons historically drove interests towards ANNs with a single hidden layer: (1) the  
274 universal function approximation theorem ([Hornik et al., 1989](#)), as it provided a compelling argument that  
275 such ANNs are fully capable of learning any function; (2) the principle of parsimony, as ANNs with fewer  
276 hidden layers are generally deemed less complex and more understandable; and (3) difficulty of training,  
277 as ANNs with more hidden layers are more complex to train (see e.g., [de Villiers and Barnard, 1993](#)).

278 So, what recently shifted the status quo towards ANNs with multiple (typically many) hidden layers?  
279 [Goodfellow et al. \(2016\)](#) attribute the beginning of this shift to the work of [Hinton et al. \(2006\)](#), where  
280 ‘unsupervised learning’ was used to pre-train deep ANNs. They show unsupervised learning could  
281 effectively initialize the network’s parameters such that the subsequent training efforts through BP would  
282 become more successful. In AI, unsupervised learning refers to a process where a model learns from  
283 ‘unlabeled’ examples, which are technically inputs with no associated output. This is as opposed to  
284 ‘supervised learning’ where examples (i.e., data points) are ‘labeled’, meaning the output associated with  
285 each input is available; this process is called ‘model calibration’ in other fields.

286 Now, one might ask how unsupervised learning can be of any help in supervised learning. A common  
287 method for this purpose uses ‘autoencoders’, which are a class of ANNs historically used for  
288 dimensionality reduction and feature learning ([Bourlard and Kamp, 1988](#)). An autoencoder is an MLP,  
289 typically trained by BP, with one (or more) hidden layer that receives input and aims to produce the same  
290 input as its output. As the number of hidden neurons in an autoencoder is smaller than the dimension of  
291 input, the input data get encoded at the hidden layer (i.e., bottleneck) with a reduced dimensionality,  
292 while preserving the information contained in the input. Autoencoders can pre-train the first layers of a  
293 deep ANN such that the weights of those layers capture the main features in input data before passing  
294 them to the next layers. After the pre-training phase by unsupervised learning, the ANN is fully trained in  
295 the conventional supervised manner, using the actual output data and algorithms such as BP.

296 While the third wave of ANN popularity began by leveraging unsupervised learning to train deep ANNs,  
297 [Goodfellow et al. \(2016\)](#) argue the interest has gradually shifted back to the classic learning algorithms,  
298 such as BP, even for training deep ANNs. Those classic learning algorithms are now believed to work quite  
299 well in the DL context, perhaps due to the emergence of unprecedented computational power. In this  
300 regard, a game changer was the introduction of graphics processing units (GPUs) to the ANN community  
301 as a powerful tool to massively parallelize and thus expedite training algorithms ([Raina et al., 2009](#)). Such  
302 computational power has enabled the development of large ANNs, in terms of both depth and width. As  
303 such, ANNs with hundreds of millions (e.g., [Devlin et al., 2018](#)) or even a trillion parameters (e.g.,  
304 [Rajbhandari et al., 2019](#)) are becoming common.

305 Such a tremendous revival of the field of ANNs might seem at first surprising to those earth and  
 306 environmental scientists who have known the field for a long time. This might be due, in part, to the fact  
 307 that ANNs developed nowadays are fundamentally similar to those developed in the 1990s. Differences,  
 308 if any in an application, are often in the details. For example, following [Glorot et al. \(2011\)](#), the tendency  
 309 now is to use the rectified linear unit (ReLU), which is an unbounded function, instead of the standard  
 310 ‘sigmoidal’ activation functions (see Eq. 1). The recent boom in data science and cyberinfrastructure and  
 311 in investments by mega companies, such as Google, in this field might explain this revival, resulting in  
 312 huge successes in image processing ([Krizhevsky et al., 2017](#)) and speech recognition ([Young et al., 2018](#)).  
 313 Perhaps recent rebranding of the field under the title of ‘deep learning’ might have been in part a  
 314 marketing strategy; as cited in [Schmidhuber \(2015a\)](#), this term was first introduced by [Dechter \(1986\)](#) to  
 315 ML and by [Aizenberg et al. \(2000\)](#) to ANNs.

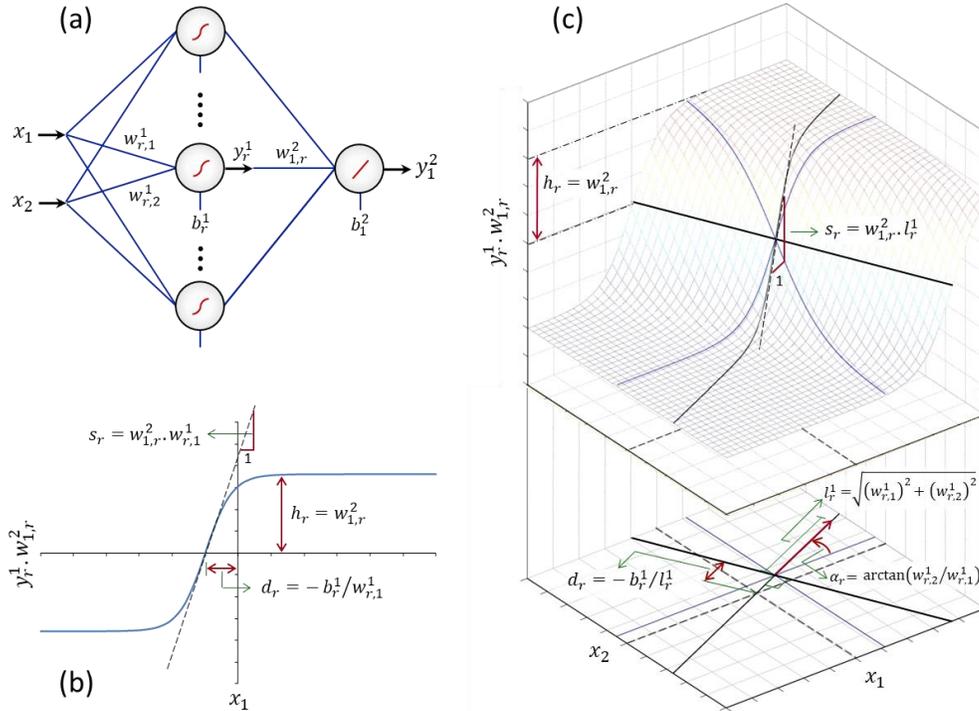
### 316 3. Geometrical Interpretation of DL

317 ANNs have always struggled with explainability and interpretability. Extensive research efforts have  
 318 endeavored to peer inside the ‘black box’ of ANNs, via various forms of sensitivity analysis (see Section  
 319 3.4 of [Razavi et al. \(2021\)](#) for a review) or geometrical or other types of interpretations (e.g., [Benítez et al., 1997](#);  
 320 [Tickle et al., 1998](#); [Castro et al., 2002](#); [Wilby et al., 2003](#); [Xiang et al., 2005](#); [Razavi and Tolson, 2011](#)).  
 321 In particular, [Razavi and Tolson \(2011\)](#) developed a geometrical interpretation of ANNs, based on  
 322 which they recast ANNs with respect to a new set of ‘explainable’ variables. This section uses that  
 323 geometrical interpretation to explain why deeper ANNs are more powerful.

#### 324 3.1. A perceptron

325 An MLP is in principle made of a number of perceptrons. Consider an MLP with a single hidden layer with  
 326 a sigmoidal activation function, as shown [Figure 2a](#). Each hidden neuron, e.g., the  $r^{\text{th}}$  neuron, is a  
 327 perceptron whose output  $y_r^1$  is multiplied by the weight  $w_{1,r}^2$  before entering the output neuron. This  
 328 hidden neuron, when only having one input  $x_1$ , forms a functional relationship such as that shown in  
 329 [Figure 2b](#). This ‘sigmoidal unit’ can be characterized by three variables: ‘slope’, ‘location’, and ‘height’.  
 330 There is one-to-one mapping between these variables and the original network variables,  $w_{r,1}^1$ ,  $b_r^1$ , and  
 331  $w_{1,r}^2$ , as shown in the figure. As such, one can directly control the shape of the sigmoidal unit through  
 332 slope, location, and height, and where needed, map them onto the network’s original variables. The  
 333 benefit of doing so is that, unlike the unintuitive original variables, the new variables are geometrically  
 334 interpretable.

335 [Figure 2c](#) shows the geometry of a perceptron with two inputs,  $x_1$  and  $x_2$ . In this case, the resulting  
 336 sigmoidal unit forms a plane that can be characterized by slope, location, and height, plus an additional  
 337 variable called ‘angle’ that specifies the direction toward which the sigmoidal unit is facing. This geometry  
 338 can be extended to perceptrons with three or more (say  $D$ ) inputs, where the sigmoidal unit becomes a  
 339 *hyperplane*, characterized by a slope, location, and height and  $D-1$  angles. Full details of this geometrical  
 340 interpretation, and how it works in practice, are available in [Razavi and Tolson \(2011\)](#). As shown in the  
 341 next section, ANNs can approximate any function by putting together a large number of such sigmoidal  
 342 units.



343

344 **Figure 2.** (a) An MLP with a sigmoidal hidden layer and linear output layer. (b) The sigmoidal line formed  
 345 by the  $r^{\text{th}}$  hidden neuron when the network has only one input,  $x_1$ . (c) The sigmoidal plane formed by the  
 346  $r^{\text{th}}$  hidden neuron when the network has two inputs,  $x_1$  and  $x_2$ . A sigmoidal line can be defined by three  
 347 variables that are related to the original weights and biases:  $h_r$  is the ‘height’ of the tails,  $s_r$  is the ‘slope’  
 348 of the tangent line at the inflection point, and  $d_r$  is the ‘location’ of the inflection point with respect to  
 349 the origin. A sigmoidal plane can be defined based on those three variables as well as  $\alpha_r$ , which is the  
 350 ‘angle’ of the normal vector perpendicular to the plane.  $l_r^1$  is the length of this vector. This geometry can  
 351 be extended to MLPs with any number of inputs (see Razavi and Tolson, 2011).

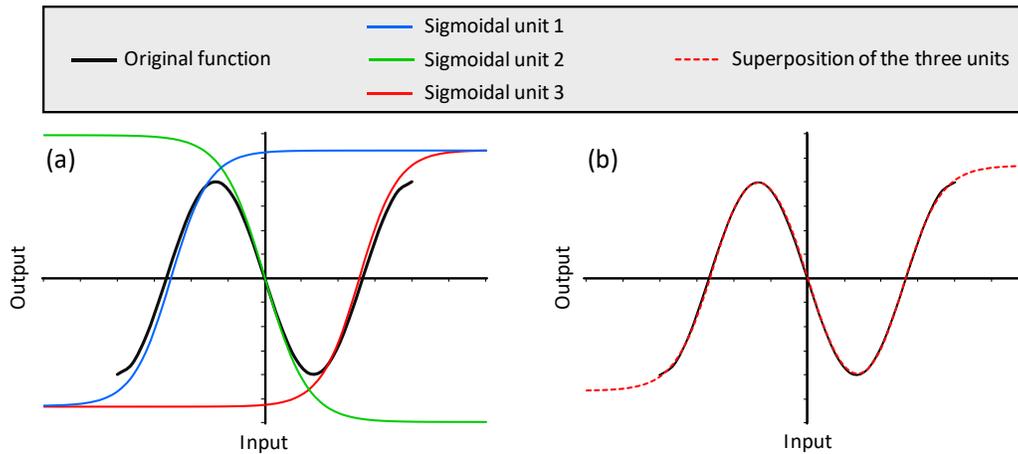
352

### 353 3.2. ANNs with one hidden layer

354 Single-hidden-layer ANNs are capable of approximating any function by combining, in parallel, as many  
 355 sigmoidal units as required. For example, suppose the underlying function to approximate is the sine  
 356 function shown in **Figure 3a**. Three sigmoidal units, with equal heights, equal absolute slopes, and  
 357 different locations, are required in parallel to represent the features of the function. These three units  
 358 can be produced by the hidden layer of an ANN and feed into a linear output layer, where they are  
 359 summed to approximate the sine function, as shown in **Figure 3b**.

360 For problems with two or more inputs, the function approximation is not as straightforward. For example,  
 361 suppose the objective in a two-input problem is to approximate the dome-like feature shown in **Figure**  
 362 **4a**. A single-hidden layer ANN with four sigmoidal hidden neurons and one linear output neuron would  
 363 approximate the dome part of the surface, as shown in **Figure 4b**. In such an ANN, four sigmoidal units  
 364 with equal heights, equal slopes, equal locations, but different angles ( $90^\circ$  apart) would be summed. The  
 365 performance of this ANN, however, is unacceptable, as it creates erroneous features on the tails.

366



367

368 **Figure 3.** (a) An original sine function and three sigmoidal units, each approximating a part of the sine  
 369 function. (b) Output of the ANN that superposes the three sigmoidal units.

370

371 **Figure 4c** shows the performance of a network with eight sigmoidal units, all having the same heights,  
 372 slopes, and locations, but different angles,  $45^\circ$  apart. With more sigmoidal units at work, the performance  
 373 at the tails is improved, producing less erroneous features. Almost 40 hidden neurons are required, as  
 374 shown in **Figure 4d**, to generate smooth tails, similar to the original function shown in **Figure 4a**. This  
 375 example provides a geometrical proof for the universal function approximation theorem of **Hornik et al.**  
 376 **(1989)** because, in principle, any function could be approximated by a combination of such dome-like (i.e.,  
 377 basis) functions. The challenge, however, is that many (possibly an excessively large number of) hidden  
 378 neurons may be required for a given problem to attain a desired level of approximation accuracy.

379

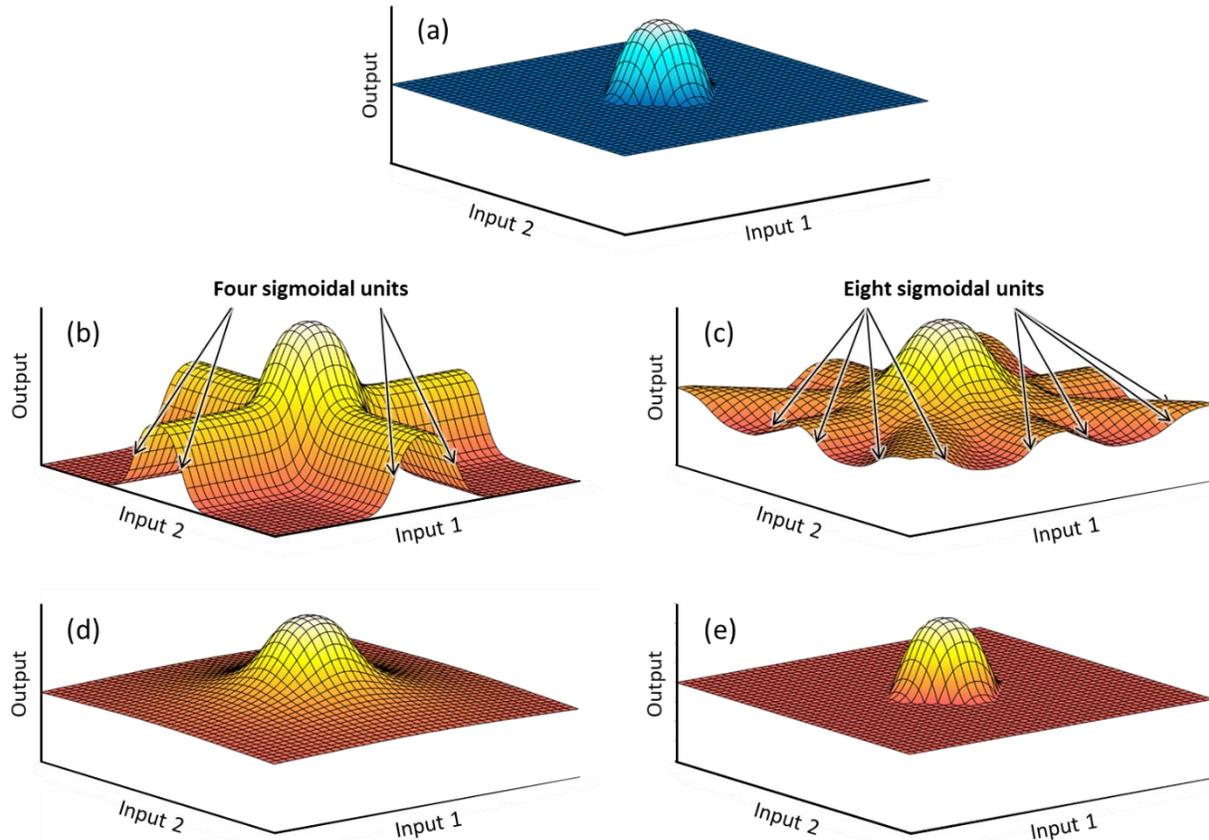
### 380 3.3. Why more than one hidden layer?

381 As proven by **Hornik et al. (1989)**, and geometrically shown in the example above, ANNs with a sigmoidal  
 382 hidden layer and a linear output layer are capable of approximating any function with any desired level of  
 383 accuracy. So, one may wonder about the need to have deeper ANNs. This section attempts to answer this  
 384 question via an example.

385 Let us look back at the original function we aimed to approximate in **Figure 4a**. Only four sigmoidal units  
 386 were required, as seen in **Figure 4b**, to reproduce the dome-like feature at the center. But can we  
 387 somehow smooth the tails? Yes, all that is needed is a second layer with a nonlinear activation function  
 388 (e.g., sigmoidal) to deactivate any feature that is under a threshold. In other words, in this process, the  
 389 geometry formed by the sigmoidal units in the first layer filters through another sigmoidal unit that  
 390 bounds that geometry. **Figure 4e** shows how adding the second non-linear layer enables the network to  
 391 reproduce the original function, with only four neurons in the first hidden layer.

392 Similar to single-hidden-layer ANNs, those with two sigmoidal hidden layers and one linear output layer  
 393 can approximate any function by putting the dome-like functions side by side. In the example shown,  
 394 however, ANNs with more than one hidden layer would require significantly fewer neurons.

395



396

397 **Figure 4.** (a) Original dome-like function. Performance of ANNs with (b) four sigmoidal hidden neurons  
 398 and a linear output neuron, (c) eight sigmoidal hidden neurons and a linear output neuron, (d) 40  
 399 sigmoidal hidden neurons and a linear output neuron, and (e) four sigmoidal hidden neurons and a  
 400 sigmoidal output neuron.

401

402 Another related consideration is that, in many problems, only a small part of the input space is active. In  
 403 other words, some combinations of the different inputs might not occur in reality and therefore the  
 404 accuracy of the model might not matter much in the regions of input space containing those  
 405 combinations. For example, consider a case similar to one shown in **Figure 4b**, where the corners on the  
 406 input space do not show up in the data available. A hydrological example is where snowfall and  
 407 temperature are two inputs to ANNs. Because snowfall would never occur along with high temperature,  
 408 the respective part of the input space always remains inactive. In such cases, single-hidden-layer ANNs  
 409 may look just as good in terms of performance as an ANN with more hidden layers. This might be one  
 410 reason why single-hidden-layer ANNs with only a limited number of hidden neurons have reportedly  
 411 worked very well.

412 In general, shallower ANNs are special cases of deeper, more flexible ANNs. However, the training of  
 413 deeper ANNs has been historically much more difficult because of the now well-known problem of  
 414 ‘vanishing and exploding’ gradients. This problem relates to the fact that the partial derivatives of a loss  
 415 function (Eq. 2) with respect to weights and biases in first layers, obtained via the chain rule of  
 416 differentiation, tend to become very small (i.e., close to zero) or very large (i.e., exponentially growing or  
 417 fluctuating). Improved algorithms along with higher computational power have now made possible the  
 418 training of very deep ANNs (Schmidhuber, 2015b).

#### 419 4. Relevance of Occam’s razor and equifinality?

##### 420 4.1. Issues with the complexity of ANNs

421 ANNs are known for their *hyper-flexibility* in fitting data, owing to their enormous degrees of freedom.  
 422 For example, consider a problem with five inputs and one output. A single-hidden-layer ANN with 10  
 423 hidden neurons would have 71 tunable parameters (60 weights and 11 biases), and adding a second 10-  
 424 neuron hidden layer would result in a network with 181 parameters (160 weights and 21 biases). Compare  
 425 that with linear or quadratic regression models for the same problem, which would have six or 21 tunable  
 426 parameters, respectively. Such large degrees of freedom, manifest in large numbers of parameters,  
 427 encountered in the field of ANNs do not seem consistent with a basic principle in statistical modelling:  
 428 *Occam’s razor*.

429 Occam’s razor, or principle of parsimony, indicates that simpler hypotheses or models should be preferred  
 430 over more complex ones. In other words, those models that serve the purpose with as few parameters as  
 431 possible should be chosen. However, many data-driven modellers, in particular in the field of ML, have  
 432 arguably abandoned Occam’s razor. For example, ANN users typically do not try simpler model types such  
 433 as regression for the problem at hand. And, when using ANNs, they do not necessarily look for the most  
 434 parsimonious network. Note that some literature proposes systematic approaches to choose a network  
 435 structure based on growing, pruning, or other strategies (e.g., Reed, 1993; Teoh et al., 2006; Xu et al.,  
 436 2006). In practice, however, such approaches have been of limited use and most ANN users choose the  
 437 network structure on an *ad hoc* basis or by trial-and-error (see a survey by Razavi et al., 2012a). Recently,  
 438 giant ANNs with hundreds of millions of parameters or more have become widespread (Devlin et al., 2018;  
 439 Rajbhandari et al., 2019).

440 In addition, *equifinality*, a common and widely discussed issue in process-based modelling (Beven and  
 441 Freer, 2001), is not generally discussed or considered an issue in the context of ANNs. Equifinality concerns  
 442 the fact that, in most cases, different model structures and parameter values can lead to identical  
 443 modelling results. In other words, model structure and parameters are not uniquely identifiable from data  
 444 (Guillaume et al., 2019). This is despite the fact that, loosely speaking, the level of equifinality of ANNs is  
 445 much larger than other types of models because of their massively parallel nature in producing model  
 446 outputs.

447 So, how does DL handle the above issues? The answer is ‘indirectly’, by trying to avoid their undesired  
 448 implications, which are *overfitting* and *lack of generalizability*. The former refers to a situation where a  
 449 model fits the noise in the data rather than the underlying function. The latter refers to a case where the  
 450 model does poorly in ‘out-of-sample prediction’, that is predicting situations unseen in the data used for  
 451 model calibration. Various techniques are available in the ANN literature to address these issues, as  
 452 outlined in the following.

453 **4.2. Leashing the hyper-flexibility of ANNs**

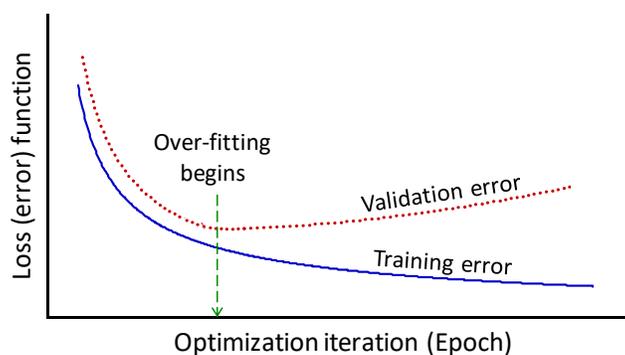
454 Techniques to control the hyper-flexibility of ANNs and to avoid overfitting fall under two general  
 455 strategies, namely ‘early stopping’ and ‘regularization’. Before reviewing these strategies in this section,  
 456 let us revisit the common data-splitting approach for calibration and validation of models.

457 ANNs and conventional mathematical models have major differences in terms of calibration and  
 458 validation. In conventional modelling practices, the available data are commonly divided into ‘calibration’  
 459 and ‘validation’ datasets. The former is used to identify the model structure and parameters, while the  
 460 latter is used to test the model performance in out-of-sample prediction.

461 In ANN practices, however, the available data are typically divided into three sets, commonly referred to  
 462 as ‘training’, ‘validation’, and ‘testing’ datasets. Any data chosen for ‘training’ and ‘testing’ in the ANN  
 463 context are respectively treated like ‘calibration’ and ‘validation’ datasets in the conventional modelling  
 464 context. The third ‘validation’ dataset is needed to leash the hyper-flexibility of the network while training.  
 465 The simultaneous use of ‘training’ and ‘validation’ datasets during training may be best described within  
 466 the ‘early stopping’ strategy, as follows.

467 The training of ANNs is an iterative process, where the network parameters are updated after each  
 468 iteration (called an ‘epoch’ in the ANN context), to minimize the loss function evaluated on the ‘training  
 469 data’ (see Eq. 2). In the ‘early stopping’ strategy, the quality of fit to the ‘validation’ dataset is also  
 470 evaluated after each epoch. Empirically speaking, as the training error decreases over time, the validation  
 471 error decreases as well for a while. But, at some particular epoch, the validation error may begin to  
 472 increase while the training error may keep decreasing (see Figure 5). This epoch is deemed to mark the  
 473 beginning of overfitting; thus, the user stops the training process. This strategy is therefore called ‘early  
 474 stopping’ in the sense that the training stops early, before it can further improve the fit to the ‘training’  
 475 dataset (for a review, see Prechelt, 1998). When the training process stops, the generalizability of the  
 476 trained network is assessed via out-of-sample prediction on the ‘testing’ dataset.

477



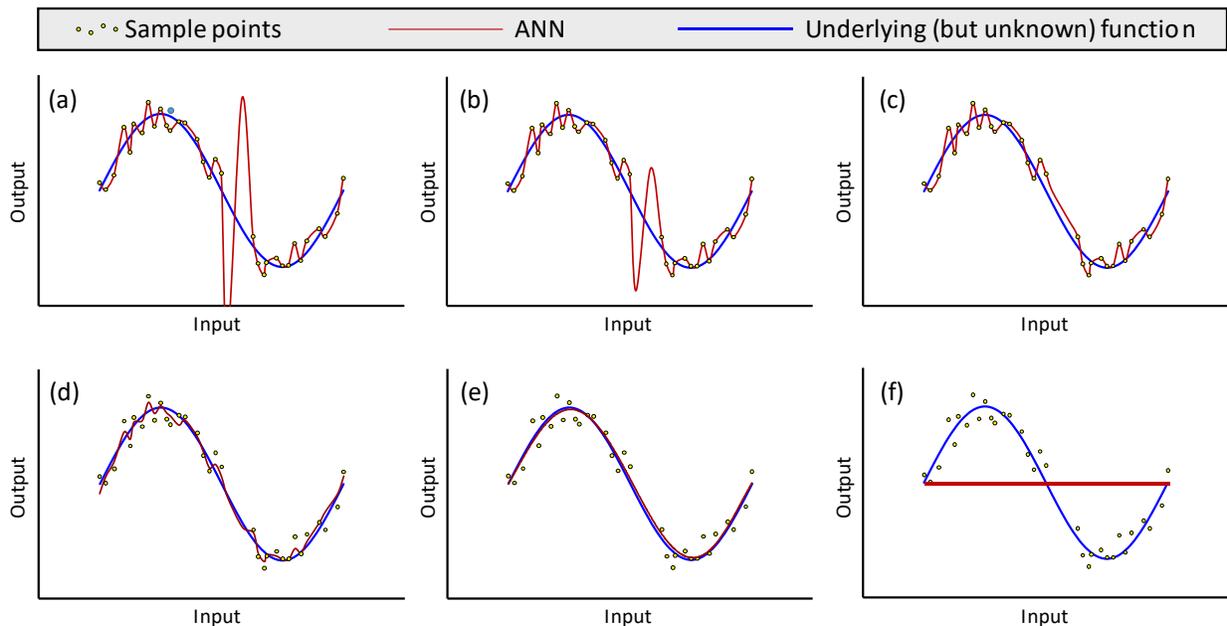
478

479 **Figure 5.** Illustration of ‘early stopping’. The loss function on the ‘training’ dataset generally decreases  
 480 with more epochs, whereas the loss function on the ‘validation’ dataset decreases early on but begins to  
 481 increase at some point, marking the commencement of overtraining.

482

483 ‘Regularization’ is another commonly used strategy to put a leash on the hyper-flexibility of ANNs. Unlike  
 484 early stopping, however, this strategy uses a ‘regularization function’ to control the ANN flexibility and  
 485 tailor it to the problem at hand. A more regularized network is one with a smoother response. A traditional  
 486 regularization function is the sum of the square of all network parameters (Krogh and Hertz, 1991), based  
 487 on the notion that, in general, the smaller the parameters of a neuron, the less activated it is. For example,  
 488 in an extreme case where all parameters of a neuron are zero, that neuron becomes fully inactive and  
 489 does not contribute a feature to the overall network response. Razavi and Tolson (2011) provide a more  
 490 efficient regularization function, based on the geometry presented in Section 3, where the regularization  
 491 function is the sum of squares of all of the slopes. This regularization function targets and removes the  
 492 unnecessary slopes, which are unsupported by data, from the overall network response.

493 But how can one balance the goodness of fit and smoothness of the network response? In practice, this  
 494 is a bi-objective optimization problem, where one objective is to minimize the error function and the other  
 495 is to minimize the regularization function. These two objective functions are commonly integrated into  
 496 one loss function via weighting schemes. Figure 6 shows how the two objectives compete in a real  
 497 example. Ideally, one may wish to achieve a performance such as that shown in Figure 6e. Doing so is not  
 498 trivial, however, because in practice the underlying function is unknown, available data are limited, and  
 499 response surfaces are multi-dimensional and cannot be easily visualized. The Bayesian regulation method  
 500 developed by MacKay (1992) and extended by Foresee and Hagan (1997) has proven useful to adaptively  
 501 assign the weights associated with each function during training.



502  
 503 **Figure 6.** Illustrative example of how regularization works to leash the hyper-flexibility of ANNs. Plot (a)  
 504 shows an extreme case with no regularization where the ANN overfits data. Plot (b) shows a case where  
 505 the regularization function is added to the loss function but marginally weighted. Plots (c) through (e)  
 506 show cases with incremental increases in the weight of the regularization function. Plot (f) shows the  
 507 other extreme case where the regularization function is dominantly weighted, making the ANN  
 508 effectively inactive. These plots are based on a real experiment, where the data sample was taken from  
 509 the underlying sine function shown and polluted with random noise.

510 A more advanced and much more complex regularization strategy is called ‘dropout’ (Hinton et al., 2012;  
511 Srivastava et al., 2014). ‘Dropout’ is a heuristic, particularly designed for deep ANNs, that randomly  
512 deactivates and then activates different neurons or groups of neurons at each epoch in the course of  
513 training. When a part of an ANN is inactivated in this process, the resulting network is called a ‘thinned’  
514 network. The ultimate prediction after training with dropout is viewed as an approximation of the average  
515 of predictions by many independent ANNs. Basically, the many different thinned networks created  
516 throughout the process are assumed to represent ANNs with different configurations and parameters.  
517 This heuristic discourages neurons to co-adapt too much and, as such, is believed to avoid overfitting.

## 518 5. Fundamental differences from other ML methods

### 519 5.1. Local versus distributed representations

520 Most ML methods, such as those based on kernel functions, are based on ‘local representations’. These  
521 methods, while forming *connectionist* networks like ANNs, represent each entity (e.g., a training sample  
522 point in the input space) via a single processing unit. For example, radial basis functions (Broomhead and  
523 Lowe, 1988), Gaussian emulator machines (Kennedy and O’Hagan, 2000), and support vector machines  
524 (Vapnik, 1998; Cherkassky and Ma, 2004) may use as many kernels as the number of training samples.  
525 Each kernel typically has a limited radius of influence in the input space, and therefore only responds to  
526 inputs located in their local neighborhood.

527 Conversely, a unique feature of ANNs is their ability to learn through ‘distributed representations’ (Hinton  
528 et al., 1986). They typically represent an entity via collective efforts distributed among multiple processing  
529 units (e.g., sigmoidal units). Unlike kernel functions, the sigmoidal units typically have large regions of  
530 influence (see e.g., Figure 2c) that overlap each other in the input space (see e.g., Figure 4b). The former  
531 figure shows that a sigmoidal unit influences the entire input space, by dividing it into three zones: lower  
532 tail, upper tail, and slope. The latter figure shows how the influences of four such sigmoidal units are  
533 superimposed to generate the network response.

### 534 5.2. Implications for users

535 The use of distributed representations has several practical implications. To the author’s knowledge, these  
536 include:

- 537 • **Transparency:** The internal functioning of methods based on local representations is more  
538 transparent. Local representations are the most straightforward and easy-to-interpret way of  
539 learning, whereas distributed representations can be complex, often leading to emergent  
540 properties that cannot be easily explained by local representations (Hinton et al., 1986).
- 541 • **Learning difficulty:** Distributed representations are more difficult and time-consuming to learn.  
542 In local representations, the role of each processing unit may be assigned independently of the  
543 other units, but in distributed representations, many processing units may be configured together  
544 in complex ways to represent a feature in the data.
- 545 • **Network size:** Distributed representations need much smaller network sizes. In general, the size  
546 of the networks based on local representations is directly proportional to the size of the dataset,  
547 in most cases with a proportionality constant of one; that is, the number of processing units

548 mirrors the number of training data samples. The size of networks based on distributed  
549 representations, however, depends on the complexity of features in the dataset, not its size.

550 • **Inexact emulation:** Networks based on distributed representations are generally ‘inexact  
551 emulators’. This means they do not exactly fit the training samples to represent the features and  
552 patterns in the data. This is unlike some other ML methods, such as radial basis functions  
553 (Broomhead and Lowe, 1988) and Gaussian emulator machines (Kennedy and O’Hagan, 2000),  
554 that are ‘exact emulators’, perfectly interpolating the training samples. Other inexact emulators  
555 include support vector machines (Vapnik, 1998; Cherkassky and Ma, 2004) and multivariate  
556 adaptive regression splines (MARS) (Friedman, 1991). Refer to Razavi et al. (2012a, Section 2.6.2)  
557 for a discussion on this issue.

558 In addition, ANNs are essentially multioutput models because they can have as many output neurons as  
559 required for a given problem. This means a single ANN can simultaneously predict different variables while  
560 accounting for their possible cross-correlations. Many other ML methods are, however, single output  
561 models. For example, in the case of support vector machines, one need to develop two independent  
562 models to be able to predict two different variables in a system. Refer to Razavi et al. (2012a, Section  
563 2.6.5) for an extensive discussion on this matter.

## 564 6. How to introduce order, time-dependency, and memory

565 MLPs provide *static* mapping from inputs to outputs. However, many applications require mappings with  
566 a formal representation of time evolution and memory. To enable MLPs to do so, two general sets of  
567 tools, and their combination, have been used in the literature: (1) tapped delay lines and (2) recurrent  
568 connections. These tools are explained in the following.

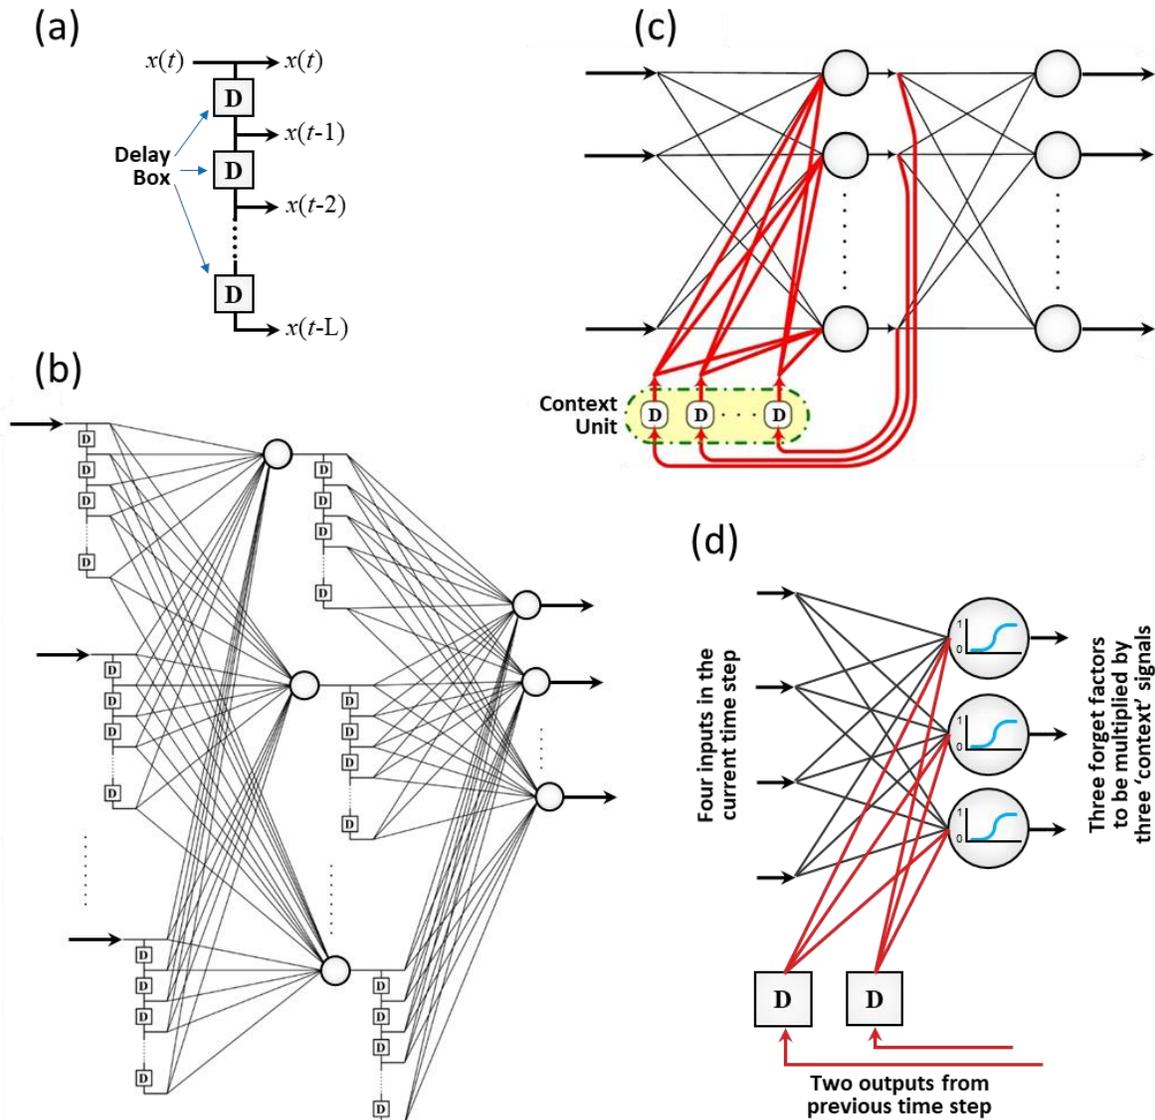
### 569 6.1. Tapped delay lines

570 A tapped delay line (TDL) consists of a certain number of time delay operators arranged in an incremental  
571 order (Figure 7a). TDLs can be installed on any internal connection weights of MLPs to represent time  
572 explicitly. The resulting ANN shown in Figure 7b, commonly referred to as a ‘time delay neural network’  
573 (TDNN; Waibel et al., 1989), has been widely used in a range of time-series processing applications. As  
574 such, TDNNs possess a *static memory with an adjustable length*. This length can be viewed as a  
575 hyperparameter to be tuned during training, along with network structural properties such as the  
576 numbers of layers and neurons in each layer.

577 Adding TDLs to an MLP significantly increases the number of tunable parameters. For example, a standard  
578 MLP with three inputs and 10 neurons in the first hidden layer would have 30 weights in that layer, while  
579 adding TDLs with a length of five to the inputs would result in an additional 50 weights (80 in total) to be  
580 trained.

581 TDNNs can be viewed as a special case of ‘convolutional neural networks’ (CNNs; Lawrence et al., 1997).  
582 CNNs, which have proven capability in image recognition, apply a ‘moving window’ approach with linear  
583 filtering to the inputs to the first and possibly other layers to preserve spatial orders in those data. TDLs  
584 essentially function in the same way as such moving windows but along one dimension only that  
585 represents time. In general, CNNs can be configured for any data with any number of dimensions.

586



587

588 **Figure 7.** (a) A tapped delay line (TDL), receiving the scalar  $x(t)$  at each time step  $t$  and outputting the  
 589 vector  $[x(t), \dots, x(t-L)]$ , where  $L$  is the length of the TDL. (b) A time delay neural network (TDNN) with one  
 590 hidden layer and TDLs installed on the input and hidden layers. (c) A recurrent neural network (RNN)  
 591 with one hidden layer and recurrent connections from the hidden neurons to themselves. In case of long  
 592 short-term memory (LSTM) networks, the context unit contains three 'gate layers' that adjust the  
 593 properties of the network's memory. (d) A gate layer of an LSTM with four inputs, two outputs, three  
 594 'context' signals that evolve through time steps.

## 595 6.2. Recurrent connections

596 TDLs, as described in [Section 6.1](#), explicitly represent time with a memory unit of limited length. Unlike  
 597 TDLs, recurrent connections, first introduced by [Jordan \(1986\)](#), enable ANNs to account for time evolution  
 598 based on an implicit memory concept, which is theoretically of unlimited length and is highly context  
 599 dependent ([Elman, 1990](#)). Recurrent connections receive the outputs of a layer at every time step and

600 feed them back to the same or some other layer in the next time step. Technically, they do so via a ‘context  
601 unit’ that stores those outputs in a set of delay boxes (**Figure 7c**). Recurrent connections can be installed  
602 on one or more layers (e.g., Jordan, 1986; Elman, 1990) or locally on some select neurons (e.g., Frasconi  
603 et al., 1992).

604 An MLP enabled with recurrent connections is commonly called a ‘recurrent neural network’ (RNN). An  
605 RNN can possess many more tunable parameters compared to an MLP with the same number of layers  
606 and neurons. Using the example given in **Section 6.1**, an MLP with three inputs and 10 neurons in the first  
607 hidden layer would have 30 weights in that layer, whereas adding recurrent connections to that layer  
608 (e.g., **Figure 7c**) would add 100 more weights (130 in total) to that layer.

609 Unlike TDNNs that possess a short-term memory, RNNs in theory can represent long-term dependencies  
610 in the input sequence as well. In practice, however, recurrent connections have difficulty representing  
611 long-term memory because they can easily get dominated by short-term memory. In other words, even  
612 very small features arising from short-term dependencies tend to mask features arising from long-term  
613 dependencies. In addition, RNNs are prone to the ‘exploding and vanishing’ gradients problem in their  
614 training (Bengio et al., 1994). This is because RNNs, even with a single hidden layer, are in principle deep  
615 networks implicitly possessing an infinite number of recursive layers.

616 To explicitly account for and balance both short- and long-term dependencies in input sequences,  
617 Hochreiter and Schmidhuber (1997) introduced a new type of RNNs, called ‘long short-term memory’  
618 (LSTM). They extended and further parametrized the ‘context’ (also called ‘cell’) such that the network  
619 can more explicitly control what information to hold over time and what to forget. The LSTM’s context  
620 unit modulates not only the outputs in the previous time step but also the inputs to the network in the  
621 current time step. It does so via three independent layers of neurons arranged in the so-called ‘forget  
622 gate’, ‘input gate’, and ‘output gate’ layers. The neurons of each ‘gate layer’ as shown in **Figure 7d**, at  
623 each time step, receive recurrent connections as well as the new input to the network, and generate their  
624 response between *zero* and *one* via using a logistic function. These responses are then multiplied by their  
625 respective signals flowing through the context, which means a value of zero would kill a signal whereas a  
626 value of one would fully preserve it. Due to the additional weights and biases in the gate layers, an LSTM  
627 typically has many more tunable parameters than a conventional RNN.

628 LSTMs are now perhaps the most popular and widely used type of ANNs with memory. However, LSTMs  
629 took a long time (more than a decade) to become known and mainstream, particularly beyond their core  
630 computer science community. Their widespread application nowadays owes to recently developed  
631 software tools such as *TensorFlow* that efficiently implement variations of LSTMs for a range of problems.

### 632 **6.3. Training considerations when the order of data matters**

633 The training of memory-enabled ANNs, such as TDNNs, RNNs, etc., is different from that of standard ANNs  
634 in terms of the way time-ordered data are presented to the network. To train standard ANNs, the data  
635 entries are typically presented randomly. In memory-enabled ANNs, however, the data entries should be  
636 presented in order of occurrence so that the structure of the time dependency is preserved. While this  
637 point might seem trivial, it requires careful attention in practical applications.

638 Another point to consider in the training of memory-enabled ANNs is that all data entries are typically  
639 viewed to have equal importance, regardless of their location in the sequence. When used in an online

640 operational forecast, however, the ‘forgetting factor’ approach can be used to discount older samples.  
641 This approach allows the network to adapt to non-stationary environments, where more recent data are  
642 more representative of the underlying processes than older data (Razavi and Araghinejad, 2009).

643 Lastly, elements of TDNNs and RNNs can be combined in a variety of ways. A well-known combination is  
644 ‘time-delay recurrent neural networks’ developed by Kim (1998) and used in various applications such as  
645 long-term precipitation forecasting in Karamouz et al. (2008); see Razavi and Karamouz (2007) for a  
646 comparison of MLP, TDNN, RNN, and TDRNN in the context of flood forecasting. While such combinations  
647 may show improved modelling power compared to other ML or statistical methods, the attribution of  
648 memory gains to the different elements can arguably be challenging, if possible at all.

## 649 **7. ML versus process-based modelling – An experiment**

650 Machine learning has been extensively used to model systems for which process-based (also called  
651 mechanistic) models are also available. Mechanistic models are based on the physics governing the  
652 underlying processes and are therefore typically evaluated based on both their physical realism and  
653 goodness of fit to data. ML, however, does not do much, if anything, with the underlying physics while  
654 reportedly doing a superior job in fitting data, even in out-of-sample prediction. A fairly large body of  
655 literature benchmarks ML techniques, particularly ANNs, against mechanistic models. Examples of such  
656 comparisons in the context of hydrologic modelling include Hsu et al. (1995), Tokar and Markus (2000),  
657 Wilby et al. (2003), Kratzert et al. (2018), and Kratzert et al. (2019). Some studies, such as Wilby et al. (2003),  
658 also detected correlations between the weights of ANNs and state variables of mechanistic models as a  
659 way to verify that ANNs can capture the underlying processes in a hydrologic system.

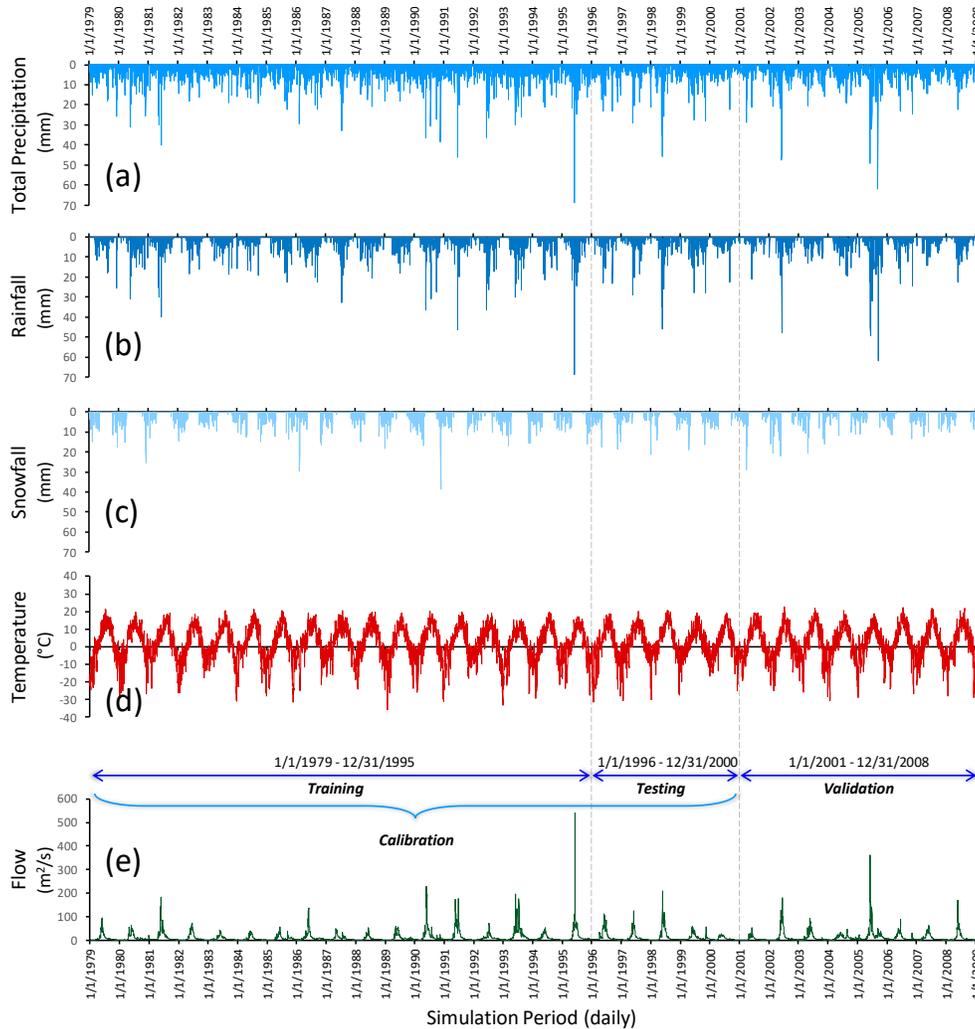
660 This section provides an experiment that runs and compares both types of models for the same problem  
661 and walks the reader through all of the steps involved. In particular, the processes around calibration and  
662 validation, role of physics, and interpretations of out-of-sample prediction are discussed. This experiment  
663 is performed in the context of hydrologic modelling, which has seen tremendous progress over the years  
664 with respect to both ML and mechanistic modelling.

### 665 **7.1. Data and models**

666 The case study used aims to model the hydrologic system of the Oldman River watershed in Alberta,  
667 Canada. This watershed has an area of 1434.73 km<sup>2</sup> at Waldron's Corner with a long-term average  
668 temperature of 2.2 °C. On average, this watershed receives 611 mm of precipitation (rainfall + snowfall)  
669 annually and generates 11.7 m<sup>3</sup>/s of river flow. **Figure 8** shows the 30-year long daily time series data  
670 used. The first 22 years were used for model ‘calibration’ (i.e., the ‘seen’ data in model development) and  
671 the last eight years for model ‘validation’ (i.e., the ‘unseen’ data in model development). The first three  
672 months of the calibration period were used for model spin-up. In the case of DL, the calibration period  
673 was further broken into ‘training’ (17 years) and ‘testing’ (5 years) periods, the latter for early stopping of  
674 the training process to avoid overfitting. Note that, as explained in **Section 4.2**, the naming convention in  
675 the ML context for the ‘validation’ and ‘testing’ periods is often the other way around.

676 To model this system, an LSTM configuration was chosen here as a state-of-the-art DL model that accounts  
677 for time dependency and memory. The inputs to the LSTM model are daily precipitation and temperature  
678 (**Figures 8a and d**) and the output is the concurrent flow (**Figure 8e**). The LSTM structure was rather  
679 arbitrarily chosen to have one hidden layer with five neurons, resulting in 166 calibration parameters. For

680 benchmarking purposes, a classic hydrologic model called HBV (Lindström et al., 1997), as implemented  
 681 in HBV-SASK (Razavi et al., 2019), was used. HBV-SASK is based on a conceptualization of physical  
 682 principles governing the water movement in a watershed using 12 calibration parameters. Each of these  
 683 parameters has a physical interpretation and a physically justified feasible range (see Figure 9 and Table  
 684 2 of Razavi et al., 2019). Full detail (including data) of this Oldman River watershed case study, which has  
 685 been developed for educational purposes, is available in Razavi et al. (2019).  
 686



687  
 688 **Figure 8.** Dataset used for the modelling experiment with ML and mechanistic modelling. (a) Measured  
 689 precipitation time series (rainfall + snowfall). (b) Estimated rainfall time series (precipitation when  
 690 temperature  $\geq 0$  °C). (c) Estimated snowfall time series (precipitation when temperature  $< 0$  °C). (d)  
 691 Measured temperature time series. (e) Measured river flow time series. The training period was used for  
 692 LSTM training, while the testing period was used for early stopping. The calibration (training + testing)  
 693 period was used for HBV calibration. The validation period was used to evaluate the performance of  
 694 both LSTM and HBV in out-of-sample prediction.

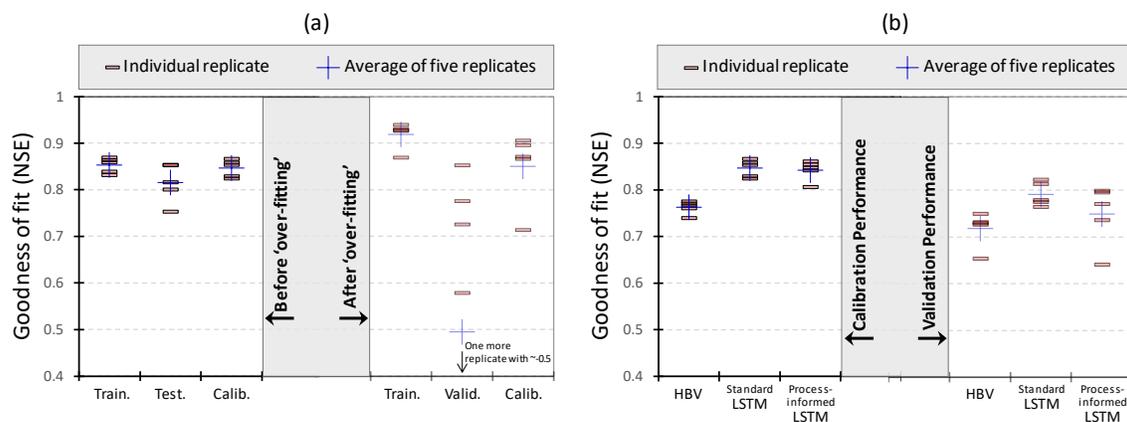
695

696 **7.2. Model performance in calibration**

697 The model calibration problem was cast as an optimization problem that tries to maximize the goodness  
 698 of fit to data by tuning the model parameters, with the Nash-Sutcliffe efficiency (NSE; Nash and Sutcliffe,  
 699 1970) as the objective function. NSE is essentially a normalized version of mean squared errors computed  
 700 as  $1 - [\text{VAR}(\text{errors})/\text{VAR}(\text{observations})]$ . As such, an NSE of one indicates a perfect fit, and an NSE of zero  
 701 indicates the model prediction is not any better than the average of observations. As a rule of thumb,  
 702 hydrologists often call an NSE of 0.7 and higher an acceptable fit.

703 The LSTM model was calibrated using BP with the early-stopping strategy to avoid overfitting. In each  
 704 epoch, the training period data were used to update the network parameters, while the testing period  
 705 data were used to detect possible overfitting. Five independent replicates of LSTM calibration (with  
 706 different initial random seeds) were conducted to account for possible variability of model performance.  
 707 **Figure 9a** shows the training results of the five replicates compared to a case where the training would  
 708 not have stopped. As expected, the LSTM performance keeps improving in training, whereas in testing it  
 709 begins to significantly degrade at some point. The objective function in training came very close to one  
 710 after many more epochs but with very poor performance in testing (not shown).

711 The HBV-SASK model was calibrated by a multi-start Newton-type optimization algorithm. Similar to  
 712 LSTM, five independent replicates of HBV-SASK calibration were run. **Figure 9b** compares the performance  
 713 of HBV-SASK with that of LSTM in calibration. At this point, only check the performance of the ‘standard’  
 714 LSTM model in calibration. The figure shows all five replicates of LSTM outperform those of HBV-SASK.  
 715 Note that the calibration performance of HBV-SASK shown herein is almost the best the author has  
 716 achieved so far for this watershed. Based on these results, the superiority of LSTM over HBV-SASK in  
 717 calibration is quite significant from a hydrologic modeling point of view. The performance of the two  
 718 models in validation is discussed in **Section 7.4**, but before that let us discuss what information the two  
 719 contained prior to calibration.



720

721 **Figure 9.** (a) The performance of LSTM in training, testing, and calibration (training + testing) periods  
 722 before and after ‘overfitting’. Training of each replicate was stopped once overtraining began at epoch  
 723 numbers ranging from 30 to 110 (left panel). Then, each replicate continued to complete 250 epochs in  
 724 total to merely evaluate the impact of overfitting (right panel). (b) A comparison of LSTM and HBV in  
 725 out-of-sample prediction. Standard LSTM and process-informed LSTM are discussed in **Sections 7.4 and**  
 726 **7.5**, respectively.

### 727 **7.3. What about *a priori* information encoded in models?**

728 At this point, let us step back and investigate what we have achieved in terms of learning from data for  
729 both the LSTM and HBV-SASK models. The development of the LSTM model was not based on any *a priori*  
730 knowledge of how a watershed system works and the governing physical principles. As such, the model  
731 learned everything from scratch merely using examples from data. Basically, the model started with a fully  
732 randomized internal configuration controlled by a large number (i.e., 186) of parameters and then tuned  
733 those parameters to adapt the internal functioning of LSTM to the underlying real-world system  
734 represented in the data. **Figure 10a** shows the LSTM performance of arbitrarily chosen replicates before  
735 and after calibration. The model response to inputs before calibration seems to be completely random  
736 but, after calibration, the model response has learned to closely follow the underlying system response.

737 Unlike LSTM, HBV-SASK encodes the expert knowledge available in the field of hydrology. This model is a  
738 collection of conservation of mass equations and process parametrizations that represent how  
739 hydrologists conceptualize the way a watershed works. This ‘physically based’ modelling structure is  
740 presumably able to emulate the behavior of any watershed by tuning only 12 parameters. **Figure 10b**  
741 shows how the model performs before calibration, with parameter values chosen to be at the midpoint  
742 of their ranges, and after calibration. The figure shows the ‘uncalibrated’ model responds reasonably to  
743 the inputs; it generally captures the timing of flows and emulates the low flow segments well but is overly  
744 responsive to large precipitation events, generating spurious spikes in flows. Calibration, either manual  
745 by expert knowledge or automatic as done here via optimization, can fix the discrepancies and fit the  
746 model output to observations.

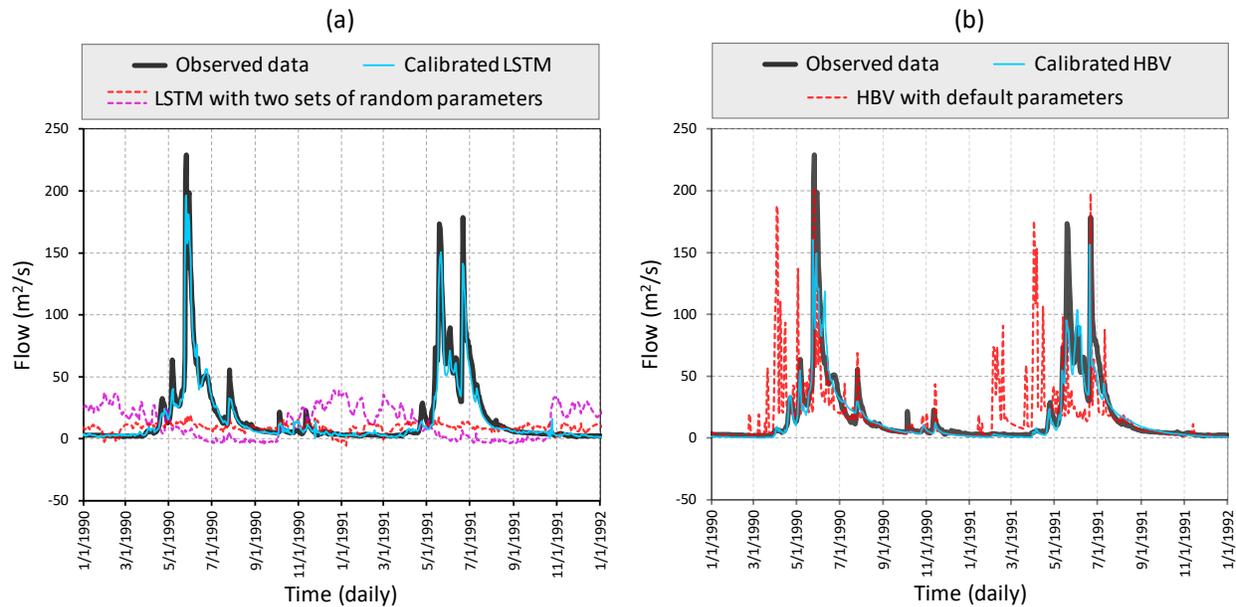
747 So, a fundamental difference between the two approaches is now clearer: using a mechanistic model is  
748 about directly using a wealth of expert knowledge available in a scientific field while using ML is about  
749 learning everything from scratch directly from data. This difference is manifest in the number of  
750 parameters that need to be tuned to achieve a reasonable performance. Notably, the LSTM model  
751 achieved a better performance in emulating observations after calibration, as evident in a comparison of  
752 **Figures 10a and b**. However, in any modelling exercise, one needs to ensure the model gives the right  
753 answer for the right reasons (Kirchner, 2006). That is why proper model evaluation in out-of-sample  
754 prediction is critically important, as discussed in the next section.

### 755 **7.4. Model validation: Standard versus true out-of-sample prediction**

756 In general, validation and verification of mathematical models are very challenging in some scientific  
757 disciplines, if possible at all (Oreskes et al., 1994). The standard practice, however, is to test the  
758 performance of the model under investigation in terms of reproducing some historical record not seen  
759 during model calibration (Klemeš, 1986a), a process called ‘out-of-sample prediction’ in this paper. **Figure**  
760 **9b** shows the results of such practice in the validation period set in **Figure 8** for both the LSTM (standard)  
761 and HBV models. In this case, both models do reasonably well from a hydrologic point of view, with LSTM  
762 outperforming HBV across all replicates. In addition, and as expected, both models produced slightly lower  
763 NSE values in validation compared to those in calibration.

764 The above so-called ‘model validation’ is inherently partial (Oreskes et al., 1994). While the performance  
765 of LSTM appears to be better than that of HBV in a ‘relative’ sense, one needs to take extra care before  
766 making such a conclusion. As argued by Klemeš (1986a) more than three decades ago, a strong  
767 assumption in this type of validation is that the conditions under which the model will be used will be

768 similar to the conditions under which the model has been developed and calibrated. It is now well-  
 769 recognized that such an assumption may not hold, as many natural systems are essentially non-stationary  
 770 (Milly et al., 2008; Razavi et al., 2015). Despite such recognition, this standard model validation practice  
 771 has arguably remained unchanged (Beven, 2018).



772  
 773 **Figure 10.** What does a model learn via calibration? Performance samples of (a) LSTM and (b) HBV  
 774 before and after calibration for a select two-year period.

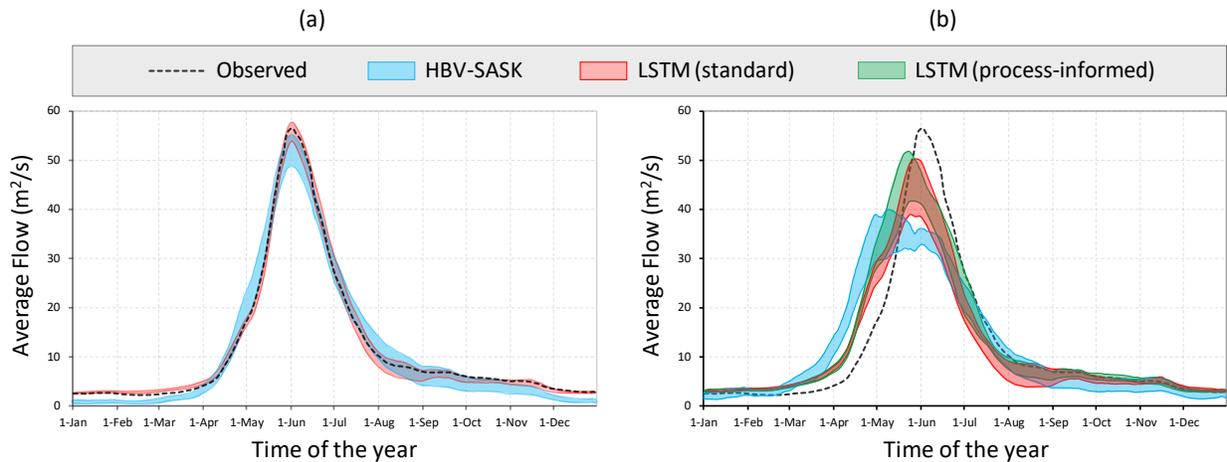
775  
 776 Here, I took a sensitivity analysis approach via a *what-if scenario* question to test and compare the  
 777 performance of both models in a ‘true’ out-of-sample prediction, basically under conditions that have not  
 778 truly been seen in the process of model development and calibration. The question is how the system  
 779 would behave if the average temperature warmed by 2 °C while everything else remained the same. To  
 780 assess this scenario, both calibrated models were fed a new temperature time series obtained by adding  
 781 2 °C to all daily temperature values of **Figure 8d**. These new inputs roughly provide a picture of what might  
 782 happen in this watershed under global warming. The modelling results under such scenarios are typically  
 783 used to inform policy making for climate change adaptation.

784 Now let us evaluate the possible changes in the watershed behavior in response to a 2 °C warming based  
 785 on the two modelling paradigms. Here, instead of looking at individual simulated time series, the possible  
 786 change in the average seasonality of flows is of interest. First, look at **Figure 11a** to check the consistency  
 787 of simulated flows for the historical period. Both models generally follow the observed seasonality, but  
 788 the range provided by the LSTM model is generally narrower and better encapsulates observations in both  
 789 low and high flows.

790 Under the new conditions, however, the two models show the two distinct behaviors shown in **Figure**  
 791 **11b**. According to LSTM, peak summer flows would decline by about 25% on average and the time of the  
 792 peak would shift backward by about a week, from the beginning of June to a time in the fourth week of  
 793 May. According to HBV-SASK, however, the changes would be more pronounced. The peak flows would

794 decline by about 35% and the flows might show two modes: the larger at the beginning of May and the  
 795 other at the beginning of June, at about the same time as the peak in the historical observations. Are such  
 796 differences not sufficiently large so as to make the user skeptical about the modelling process?

797



798

799 **Figure 11.** Long-term average daily flows throughout the year under (a) historical and (b) hypothetical  
 800 conditions. The envelopes represent the daily ranges of flows obtained by the five replicates of each  
 801 model. The curves were smoothed by a 20-year moving average filter.

802

### 803 7.5. Injecting some physics into ML

804 At this point, one may wonder about the possibility of ensuring that DL results be physically consistent,  
 805 particularly under new conditions. Let us give it a try by recasting the modelling problem based on some  
 806 understanding of the governing physics in hydrology. For example, physics tells us that the freezing point  
 807 of water is around 0 °C and, therefore, this threshold could be used as an approximation to differentiate  
 808 rainfall from snowfall on a daily basis, i.e., if the temperature on a day is above/below 0 °C, the  
 809 precipitation on that day, if any, is considered to be rainfall/snowfall (see **Figures 8b and c**). This  
 810 differentiation is actually a part of process parameterization in HBV, similar to many other hydrologic  
 811 models, via a parameter called ‘temperature threshold’ (TT) for melting/freezing and separating rain and  
 812 snow, with a feasible range from –4 to +4 °C (see **Razavi et al., 2019** for details). The warming of a  
 813 watershed would naturally change the rainfall to snowfall ratio, and so building this domain knowledge  
 814 into the LSTM model makes sense.

815 Perhaps the most straightforward way of introducing the TT concept into LSTM is via pre-processing of  
 816 the inputs. Therefore, a new LSTM model was developed and calibrated, called ‘process-informed LSTM’  
 817 in this paper, with three inputs: rainfall, snowfall, and temperature as shown in **Figures 8b, c, and d**.  
 818 Similar to the original, the new LSTM model has one hidden layer with five neurons, resulting in 186  
 819 calibration parameters. The procedure for the calibration and validation of the process-based LSTM was  
 820 the same as for the ‘standard LSTM’, already explained in **Sections 7.2 and 7.4**. **Figure 9b** compares the  
 821 performance of the process-informed LSTM with HBV and the standard LSTM. The figure shows the two

822 LSTM models perform comparably well. Process-informed LSTM results in a slightly lower average NSE in  
823 validation but, with only five replicates, this small difference should be interpreted with caution.

824 **Figure 11b** demonstrates the performance of the process-informed LSTM model in the true out-of-sample  
825 prediction. According to this model, the summer peak flows would decline by 20% on average and the  
826 time of peak would appear about two weeks earlier than in the historical record, in the third week of May.  
827 The process-informed LSTM model generated rising and falling limbs that are more consistent with those  
828 of HBV-SASK. Overall, however, the results of HBV-SASK under the new conditions are still quite different.

## 829 **7.6. So, what model should we trust: the ML or physically based model?**

830 Now the question is which one of the three models produced the most credible picture of possible  
831 watershed behavior under the new conditions. In practice, this question is very difficult to answer, if  
832 possible at all. In general, the prediction of such changes can be debated and might vary from one study  
833 to another, depending on the models and data used. Perhaps, a definite answer would need to wait until  
834 the future has come and shown such possible changes. And, from a bigger-picture point of view, models  
835 of natural systems cannot be verified or validated in true out-of-sample prediction, because those systems  
836 are never closed and not everything can be represented in a model, as argued by [Oreskes et al. \(1994\)](#)  
837 nearly three decades ago.

838 But, as scientists, we have our own perceptions and intuitions. These might be biased but still useful to  
839 provide a ground for building confidence in the credibility of a model. In the context of the case study  
840 given, previous research on the Canadian Rocky Mountains has indicated that warming *alone* will result  
841 in a considerable reduction in flows and earlier peaks in watersheds similar to the Oldman River  
842 watershed. A synthesis of research efforts under the Changing Cold Regions Network (CCRN; [DeBeer et al., 2020](#))  
843 on the cold interior of western Canada indicates a shift in timing of the spring hydrograph rise  
844 and peak flows of nearly two weeks earlier by mid-21<sup>st</sup> century, and as much as one month by the late  
845 21<sup>st</sup>-century.

846 The addition of some physics to the LSTM model in **Section 7.5** should at least intuitively improve trust in  
847 modelling results. What is worrisome is the large divergence in behavior between models that produce  
848 comparable results in standard out-of-sample prediction. This requires a more in-depth understanding  
849 and appreciation of the value of domain knowledge, as discussed in the next section.

## 850 **8. Discussion**

### 851 **8.1. What is the typically ignored value of domain knowledge in DL?**

852 True out-of-sample prediction is nothing but ‘extrapolation’ beyond the observed data and behaviors used  
853 in model development and calibration. Extrapolation is a reality that many predictive models nowadays  
854 must face because of ‘non-stationarity’ in climate and the environment ([Milly et al., 2008](#); [Razavi et al., 2015](#)).  
855 Any purely regression-type model, including those arising from DL, would be disadvantaged in  
856 extrapolation as, by definition, extrapolating would require working in parts of the problem space for  
857 which they have not received any information. Conversely, mechanistic models may be salvaged in  
858 extrapolation by the domain knowledge encoded within them.

859 But what does domain knowledge offer when it comes to extrapolation? The answer is the set of principles  
860 modulated via conservation laws (e.g., mass, energy, and momentum) and process parametrizations,

861 which represent our perceptions of how two or more variables might be related (Gupta et al., 2012). Such  
862 principles have been developed and evolved over time based on extensive observation and research by  
863 scientists and practitioners. The limits of validity of such principles are typically known. In the following,  
864 the importance of taking advantage of those principles in modelling and prediction is discussed with  
865 respect to three aspects: conservation laws, monotonicity and rates, and feedback mechanisms.

866 **Conservation laws:** In physics, a conservation law states that a specific measurable property does not  
867 change within an isolated system with time. Such a law is usually expressed as a ‘continuity equation’;  
868 that is, a differential equation equates the rate of change in storage within a control volume with the  
869 difference between what comes in and what goes out of the control volume. In hydrology, for example,  
870 continuity equations are built into mechanistic models to ensure water balance is preserved in simulations  
871 over time. ML models, however, do not automatically account for such laws and, as a result, water can be  
872 *falsely* introduced or lost in the course of simulation.

873 **Monotonicity and rates:** The knowledge base includes the general characteristics of some causal  
874 relationships between various physical variables. For example, we know from basic thermodynamics that  
875 the relationship between melt rate and available heat is *monotonic*; that is, more heat causes a higher  
876 melt rate. Furthermore, we have some rough estimate of the feasible range of the *rate* of change in one  
877 with respect to the other. Similarly, from basic hydrology we know the causal relationships governing the  
878 way a hillslope stores and releases water are generally such that a positive correlation exists between  
879 water available in the soil and its contribution to flows; more water means more flows due to gravitational  
880 forces.

881 Mechanistic models directly account for such knowledge on casual relationships. This knowledge is  
882 encoded in process parametrizations typically in the form of deterministic, monotonic functions, or rarely  
883 in hysteretic forms, with a limited number of parameters to be calibrated to the specific case study in  
884 hand (Gharari and Razavi, 2018). However, in the case of hyper-flexible models such as ANNs, such  
885 functions need to be entirely derived from data, all from scratch, and ignoring the knowledge base related  
886 to those monotonic relationships. Therefore, extrapolation runs the risk that such relationships become  
887 non-monotonic and/or have unrealistic rates, producing erroneous behaviors. This risk is exacerbated by  
888 the fact that identifying and diagnosing such errors are very difficult, if possible at all.

889 **Feedback mechanisms:** A real-world physical system is a combination of variables that interact over time,  
890 typically via a range of feedback mechanisms. Such feedback mechanisms control the internal dynamics  
891 of the system and are key to its evolution over time. For example, consider a coupled water-vegetation  
892 system in which precipitation, available soil moisture, and plant biomass interact in complex time-  
893 dependent ways, even at times creating positive feedbacks that destabilize the system’s behavior  
894 (Rodriguez-Iturbe et al., 1991; Scheffer et al., 2001). The knowledge base available about these feedback  
895 mechanisms is often built into mechanistic models, using differential equations (ordinary or partial) to  
896 describe the system dynamics. The representation of such dynamics in the making of models is important,  
897 particularly for long-term predictions and over long time scales.

898 DL models are often unable to account explicitly, perhaps even implicitly, for such long-term dynamics. If  
899 a particular dynamical behavior is present in training data, then DL can capture that behavior in its  
900 mapping from input onto output. But DL has no explicit mechanism to represent that dynamic under  
901 perturbed conditions beyond what has been recorded in the training data.

902 The bottom line is that mechanistic models are generally expected to be less prone to generating spurious  
903 behaviors in true out-of-sample prediction. Therefore, many domain experts may be inclined to trust  
904 physically based models as their behavior is constrained by physical laws that are perceived as unchanging  
905 with time. The points made in this section will become clearer in the next section, where the essential  
906 differences between DL and mechanistic modelling are discussed.

## 907 **8.2. Why is DL essentially different from process-based modelling?**

908 In the author's view, the first principles of ANNs are rooted in *connectionism*, *hyper-flexibility*, and  
909 *vigorous optimization*. These characteristics are fundamentally different from the guiding principles of  
910 developing and calibrating mechanistic models, as described in the following:

911 • *Connectionism* is an approach that orchestrates a set of simple algebraic operations in a massively  
912 parallel manner to create a model that is able to carry out complicated tasks. Following this  
913 approach, ANNs represent the response of a system under consideration to an input by summing  
914 the collective efforts of many neurons, whose roles cannot be easily attributed to individual  
915 processes involved in that system. This is unlike mechanistic modelling where each part of a model  
916 is designed to be responsible for a specific process.

917 • *Hyper-flexibility* is a characteristic of a model with excessive degrees of freedom, which can literally  
918 fit any dataset, and is not constrained by the many assumptions held by typical statistical models.  
919 ANNs are known to be hyper-flexible. Mechanistic models, however, have limited degrees of  
920 freedom depending on the knowledge base available about the processes being modelled. In  
921 general, mechanistic models tend to have just as many degrees of freedom as can be supported and  
922 constrained by available knowledge and data.

923 • *Vigorous optimization* here refers to the practice of manipulating model parameters at any cost to  
924 maximize the goodness-of-fit to calibration data. The training of ANNs is all about minimizing an error  
925 function; that is, among two competing ANNs, the one producing smaller errors in calibration and  
926 validation is the winner. Optimization is also often an essential part of mechanistic modelling to  
927 calibrate model parameters. However, in mechanistic modelling, minimizing the errors is not the  
928 goal but a means to improve the realism of the model. In other words, unlike ANNs, physical  
929 feasibility of a parameter, its identifiability, and equifinality are key considerations in mechanistic  
930 modelling.

931 The recognition of these fundamental differences is critically important when one aims to choose the  
932 correct modelling paradigm for a purpose, compare the two paradigms in a case study, or attempt to  
933 bridge the two paradigms, possibly for improved modelling performance. The following section outlines  
934 the status quo for bridging the two paradigms and some emerging trends.

## 935 **8.3. How can we bridge DL and process-based modelling?**

936 The history of research on reconciling and bridging ANNs with mechanistic modelling dates back to the  
937 early 2000s or perhaps earlier. These efforts have generally had the objective of simultaneously leveraging  
938 the strengths of the two modelling paradigms to further our knowledge and predictive ability. **Abraham**  
939 **et al. (2012)** reviewed such research in the context of hydrology and refer to it as 'hybridization'. They  
940 introduced three possible approaches for this purpose, which herein are referred to as 'surrogate  
941 modelling', 'one-way coupling', and 'modular coupling'. Seven years later, **Reichstein et al. (2019)** in an

942 influential article in Nature re-introduced and proposed the notion of ‘hybrid modelling’ and the above  
 943 three approaches as the next steps in earth science. In the following, these three approaches are  
 944 explained, and then more modern existing approaches arising from research fields beyond earth and  
 945 environmental sciences are discussed.

946 **Surrogate modelling**, alternatively called metamodeling or model emulation, refers to the process of  
 947 developing and applying a simpler, cheap-to-run model in lieu of a more complex, computationally  
 948 intensive model (Razavi et al., 2012a). In this process, a data-driven surrogate, such as an ANN, is trained  
 949 on samples of a limited number of original model runs to approximate the model response surface. The  
 950 developed surrogate model can then be used in different frameworks in conjunction with the original  
 951 model, as reviewed in Razavi et al. (2012a), in multi-query applications such as optimization and  
 952 uncertainty quantification. Example applications of ANNs as surrogates of mechanistic models include  
 953 Johnson and Rogers (2000), Behzadian et al. (2009), and Razavi et al. (2012b).

954 **One-way coupling** refers to the process combining a mechanistic model with an ML model such that the  
 955 output of the former feeds into the latter as input. A general rationale for such a combination is that a  
 956 mechanistic model may not be able to fully explain the observed data and, therefore, an ML model could  
 957 be of help in extracting any information left in the residuals of the mechanistic model. For example,  
 958 consider a case where a mechanistic hydrologic model is used for streamflow forecasting and, as  
 959 expected, some errors in model outputs are present. An ANN can be used to model such errors over a  
 960 historical period to provide some predictive ability on the errors for a time step into the future. Then,  
 961 running these two models in sequence may provide higher forecasting skills. Example applications of such  
 962 one-way coupling include Shamseldin and O’Connor (2001) and Anctil et al. (2003).

963 **Modular coupling** refers to cases where an ML model is used as a module/sub-model of a larger  
 964 mechanistic model or *vice versa*. The rationale for this type of coupling may be that a particular model  
 965 might have proven skills in representing a particular process and is therefore preferred, while other  
 966 processes are better represented by another model. Hydrologic examples are the work of Chen and  
 967 Adams (2006) and Corzo et al. (2009), in which ANNs are used as the routing module within a distributed  
 968 hydrological model. Another example is the work of Chua and Wong (2010) in which an ANN-based  
 969 hydrologic model takes the output of a kinematic wave model as one of its inputs. And, a recent example  
 970 is the work of Bennett and Nijssen (2020), in which a DL-based model for the simulation of turbulent heat  
 971 fluxes is built into a process-based hydrologic model.

972 Beyond the earth and environmental sciences community, the notion of bridging the knowledge base and  
 973 ML has a long history (e.g., see the ‘knowledge-based artificial neural networks’ by Towell and Shavlik  
 974 (1994)), but it has received significantly more attention recently. Different approaches mostly arising from  
 975 mathematics and computer science have been proposed under titles such as ‘theory-guided data science’  
 976 (Karpatne et al., 2017), ‘informed machine learning’ (von Rueden et al., 2019), and ‘physics-informed  
 977 neural networks’ (Raissi et al., 2019). Providing a full coverage of such approaches is well beyond the  
 978 scope of this paper, and many of them have been developed for specific application areas with limited  
 979 relevance to earth and environmental problems. Instead, in the following, I try to be selective and explain  
 980 three approaches that I found most relevant.

981 **Regularizing ANNs via knowledge-based loss terms**. A new regularization function can be developed  
 982 based on the available knowledge surrounding a given problem and be added to the loss function used in  
 983 training. For example, any violation of the conservation laws or monotonicity of relationships, as described

984 in **Section 8.1**, can be quantified and penalized during training. Refer to **Stewart and Ermon (2017)** for an  
 985 example application of this approach in the context of image processing.

986 **Using mechanistic model runs to augment ANN training data.** A mechanistic model can be used to  
 987 simulate the system under investigation under a range of conditions to generate ‘synthetic data’ to  
 988 augment the available training data. This approach may be particularly useful in guiding ANNs in  
 989 extrapolation beyond conditions seen in the original training data (see the discussion in **Section 8.1**). This  
 990 approach is based on the assumption that the mechanistic model used is sufficiently accurate—an  
 991 assumption that needs to be treated with caution. For an example of this approach in the field of systems  
 992 biology, see **Deist et al. (2019)**.

993 **Integrating differential equations into ANNs.** This approach is a very recent and perhaps the most  
 994 mathematically elaborate in terms of integrating the knowledge base into ANNs, primarily developed by  
 995 **Raissi et al. (2019)**. It parametrizes the known differential equations describing a system and integrates  
 996 them into the body of ANNs. The integrated model is then trained to the available data, simultaneously  
 997 inferring the parameters of the differential equations and network weights. This approach still seems  
 998 embryonic but perhaps with great potential for scientific breakthroughs.

#### 999 **8.4. What can we learn from prominent DL applications?**

1000 As outlined in **Section 1**, DL has already been used across a wide range of disciplines and applications with  
 1001 varying degrees of success. Here, and for context, consider two special and well-known cases of DL  
 1002 applications: playing chess and predicting the stock market. DL has achieved incredible, superhuman-level  
 1003 performance in chess and similar games (**Silver et al., 2018**), while its performance in stock market  
 1004 prediction has been criticized despite its widespread application (e.g., **Pearlstein, 2018**). These opposing  
 1005 outcomes may be explained as follows:

- 1006 • Chess does not possess any properties of ‘complex systems’ (**Bar-Yam, 1997**), whereas financial  
 1007 systems are essentially *complex*, with a wide range of agents interacting at a wide range of scales,  
 1008 giving rise to emergent behaviors and even black swans. Any AI-based financial services themselves  
 1009 would also be an agent influencing the stock market, even possibly inducing vicious cycles.
- 1010 • Chess can be viewed as a closed system, as no exogenous factors influence any properties or  
 1011 dynamics of the board and players, whereas stock markets are open systems and, for any analyses,  
 1012 the assumed boundary conditions depend on the analyst’s judgement.
- 1013 • Chess is a *fully observable* system, as the entire board, pieces, rules, and moves are seen by the  
 1014 players, but stock markets are only *partially observable* and some controlling elements in the market  
 1015 might be hidden to the analysts.
- 1016 • Chess is *stationary*, as the properties and governing rules of the game remain constant over time,  
 1017 whereas stock markets are *non-stationary* and their long-term dynamics and behaviors may change  
 1018 in unpredictable ways driven by political, social, economic, or natural events.

1019 So what? Earth and environmental systems arguably fall somewhere in between these two specific  
 1020 applications with respect to their four fundamental and inter-related characteristics: such systems are  
 1021 *complex, open, partially observable, and non-stationary*. Loosely speaking, understanding and predicting  
 1022 earth and environmental systems face similar challenges to those of the stock markets in terms of those

1023 four characteristics. However, unlike stock market systems that are conceived to be partially predictable  
1024 at best (Fama, 1970; Malkiel, 2003), the behaviors of earth and environmental systems are generally  
1025 believed to be predictable, with limits of predictability that have been improving as more knowledge and  
1026 data become available.

1027 The comparisons above try to convey two points. First, the revolutionary success of DL in one field of  
1028 application cannot necessarily be extended to another field of application. The context matters, and  
1029 success depends on the characteristics of the problem at hand. Second, different disciplines may cross-  
1030 fertilize DL applications and learn from one another. However, this requires more direct communications  
1031 between experts in different disciplines about common issues, which is non-trivial.

## 1032 9. Concluding remarks

1033 Deep learning has perhaps by now served every researcher and practitioner in earth and environmental  
1034 sciences communities in tasks such as image and language processing, at least through their smart phones.  
1035 Such astonishing and within-reach technologies have boosted interest in DL, and in AI in general, within  
1036 these communities, evidenced by the significant growth in the number of their research papers on DL.  
1037 Many believe the combination of AI with unprecedented data sources and increased computational  
1038 power will offer exciting new opportunities for expanding our knowledge about various earth and  
1039 environmental systems. Unsurprisingly, similar to many other innovations, AI and particularly DL  
1040 techniques are facing different views towards their future; for example, in the hydrology context Nearing  
1041 et al. (2020) suggest a DL-informed divorce from some of the current hydrological theories while Beven  
1042 (2020) advocates for the fundamental needs of a knowledge base in DL interpretation.

1043 It is certainly an exciting time for earth and environmental sciences to benefit from DL tools. We need,  
1044 however, to be mindful of any possible risk of over-excitement about the new potential and over-sellings  
1045 about the available tools. Arguably, DL in earth and environmental sciences has primarily focused on off-  
1046 the-shelf applications of methods largely developed by mathematicians and computer scientists to  
1047 problems in a new domain with no or limited considerations of the available domain's knowledge base.  
1048 The immediate risk of such practices is that the popularity of AI tools in earth and environmental sciences  
1049 would then follow the ups and downs of these tools in the areas from which they originate and the  
1050 software developed for those purposes. There is also a greater risk, in the author's view, as follows.

1051 Let us flash back to more than three decades ago, when the prominent statistician George Box (1976, p.  
1052 797-798) warned about the "mathematistry" trap, "characterized by development of theory for theory's  
1053 sake, which since it seldom touches down with practice, has a tendency to redefine the problem rather  
1054 than solve it". He argued that "there is unhappy evidence that mathematistry is not harmless. In such  
1055 areas as sociology, psychology, education, and even, I sadly say, engineering, investigators who are not  
1056 themselves statisticians sometimes take mathematistry seriously. Overawed by what they do not  
1057 understand, they mistakenly distrust their own common sense and adopt inappropriate procedures  
1058 devised by mathematicians with no scientific experience." This sentiment was then echoed by the  
1059 prominent hydrologist Vit Klemeš (1986b, p. 177 and p. 185), who said "The danger increases with the  
1060 proliferation of computerized "hydrologic" models whose cheaply arranged ability to fit data is presented  
1061 as proof of their soundness and as a justification for using them for user-attractive but hydrologically  
1062 indefensible extrapolations." He continued, "The danger to hydrology from extrapolations based on  
1063 mathematistry is that they lead it on the path of bad science."

1064 The point here is that the risk of mathematistry seems to be just as fresh as it must have been back then,  
1065 particularly when it comes to the application of AI tools in earth and environmental sciences. Due to the  
1066 very nature of such tools, this risk may even well extend to their original areas of application, to a point  
1067 that such practice has been referred to as a form of modern “alchemy”; see [Rahimi and Recht \(2017\)](#) for  
1068 the sentiment, [LeCun \(2017\)](#) for a rebuttal, and [Hutson \(2018\)](#) for a summary. This point is not to  
1069 undermine the benefits of AI technology, particularly for earth and environmental applications. Instead,  
1070 it calls for improved rigor and better appreciation of the knowledge base available. After all, it has been  
1071 long known in environmental sciences that complex models can be made to produce virtually any desired  
1072 behavior given their large degrees of freedom, as articulated by [Hornberger and Spear \(1981\)](#) three  
1073 decades ago.

1074 Having such risks in mind, the new potential afforded by AI for earth and environmental sciences is great.  
1075 To realize this potential, we need to reconcile data-driven AI techniques and the theory-driven knowledge  
1076 base. The knowledge base is at the heart of ‘traditional programming’, which is still a major building block  
1077 of process-based or mechanistic modelling in earth and environmental sciences. Clearly, the traditional,  
1078 knowledge-based programming and AI are made up of two fundamentally different world views for  
1079 problem solving and, therefore, their reconciliation will not be straightforward. This paper tried to address  
1080 some critical questions in this regard and provide some perspective for this important endeavor, in  
1081 anticipation of new breakthroughs in earth and environmental sciences in an age of big data and  
1082 computational power.

### 1083 **Acknowledgements**

1084 The archiving of the dataset used for hydrologic modelling is underway in GitHub. In the meantime, and  
1085 for the review purposes, the dataset is temporarily uploaded as Supporting Information.

### 1086 **References**

- 1087 Abrahart, R. J., Anctil, F., Coulibaly, P., Dawson, C. W., Mount, N. J., See, L. M., ... & Wilby, R. L. (2012).  
1088 Two decades of anarchy? Emerging themes and outstanding challenges for neural network river  
1089 forecasting. *Progress in Physical Geography*, 36(4), 480-513.
- 1090 Aizenberg, I., Aizenberg, N. N., & Vandewalle, J. P. L. (2000). *Multi-Valued and Universal Binary Neurons:  
1091 Theory, Learning and Applications*. Springer Science & Business Media.
- 1092 Anctil, F., Perrin, C., & Andréassian, V. (2003). ANN output updating of lumped conceptual rainfall/runoff  
1093 forecasting models, *Journal of the American Water Resources Association*, 39(5), 1269-1279.
- 1094 Ashouri, H., Hsu, K. L., Sorooshian, S., Braithwaite, D. K., Knapp, K. R., Cecil, L. D., ... & Prat, O. P. (2015).  
1095 PERSIANN-CDR: Daily precipitation climate data record from multisatellite observations for  
1096 hydrological and climate studies. *Bulletin of the American Meteorological Society*, 96(1), 69-83.
- 1097 Badran, F., Thiria, S., & Crepon, M. (1991). Wind ambiguity removal by the use of neural network  
1098 techniques. *Journal of Geophysical Research: Oceans*, 96(C11), 20521-20529.
- 1099 Bankert, R. L. (1994). Cloud classification of AVHRR imagery in maritime regions using a probabilistic neural  
1100 network. *Journal of Applied Meteorology*, 33(8), 909-918.
- 1101 Bar-Yam, Y. (1997). *Dynamics of Complex Systems*. Perseus Books, USA.

- 1102 Behzadian, K., Kapelan, Z., Savic, D., & Ardeshir, A. (2009). Stochastic sampling design using a multi-  
1103 objective genetic algorithm and adaptive neural networks. *Environmental Modelling & Software*,  
1104 24(4), 530-541.
- 1105 Benediktsson, J. A., Swain, P. H., & Ersoy, O. K. (1990). Neural network approaches versus statistical  
1106 methods in classification of multisource remote sensing data. *IEEE Transactions on Geoscience and*  
1107 *Remote Sensing*, 28(4).
- 1108 Bennett, A., & Nijssen, B. (2020). Deep learned process parameterizations provide better representations  
1109 of turbulent heat fluxes in hydrologic models. *Earth and Space Science Open Archive ESSOAr*.
- 1110 Bengio, Y., Simard, P., & Frasconi, P. (1994). Learning long-term dependencies with gradient descent is  
1111 difficult. *IEEE Transactions on Neural Networks*, 5(2), 157-166.
- 1112 Benítez, J. M., Castro, J. L., & Requena, I. (1997). Are artificial neural networks black boxes? *IEEE*  
1113 *Transactions on Neural Networks*, 8(5), 1156-1164.
- 1114 Bergen, K. J., Johnson, P. A., Maarten, V., & Beroza, G. C. (2019). Machine learning for data-driven  
1115 discovery in solid Earth geoscience. *Science*, 363(6433).
- 1116 Beven, K. (2020). Deep learning, hydrological processes and the uniqueness of place. *Hydrological*  
1117 *Processes*, 34, 3608–3613, <https://doi.org/10.1002/hyp.13805>.
- 1118 Beven, K. J. (2018). On hypothesis testing in hydrology: why falsification of models is still a really good  
1119 idea. *Wiley Interdisciplinary Reviews: Water*, 5(3), e1278.
- 1120 Beven, K., & Freer, J. (2001). Equifinality, data assimilation, and uncertainty estimation in mechanistic  
1121 modelling of complex environmental systems using the GLUE methodology. *Journal of Hydrology*,  
1122 249(1-4), 11-29.
- 1123 Bourlard, H., & Kamp, Y. (1988). Auto-association by multilayer perceptrons and singular value  
1124 decomposition. *Biological Cybernetics*, 59(4-5), 291-294.
- 1125 Box, G. E. (1976). Science and statistics. *Journal of the American Statistical Association*, 71(356), 791-799.
- 1126 Boznar, M., Lesjak, M., & Mlakar, P. (1993). A neural network-based method for short-term predictions of  
1127 ambient SO<sub>2</sub> concentrations in highly polluted industrial areas of complex terrain. *Atmospheric*  
1128 *Environment. Part B. Urban Atmosphere*, 27(2), 221-230.
- 1129 Broomhead, D. S., & Lowe, D. (1988). Multivariable functional interpolation and adaptive networks.  
1130 *Complex Systems*, 2, 321–355.
- 1131 Cabrera-Mercader, C. R., & Staelin, D. H. (1995). Passive microwave relative humidity retrievals using  
1132 feedforward neural networks. *IEEE Transactions on Geoscience and Remote Sensing*, 33(6), 1324-  
1133 1328.
- 1134 Castro, J. L., Mantas, C. J., & Benítez, J. M. (2002). Interpretation of artificial neural networks by means of  
1135 fuzzy rules. *IEEE Transactions on Neural Networks*, 13(1), 101-116.
- 1136 Chen, J., & Adams, B. J. (2006). Integration of artificial neural networks with conceptual models in rainfall-  
1137 runoff modeling. *Journal of Hydrology*, 318(1-4), 232-249.

- 1138 Cherkassky, V., & Ma, Y. Q. (2004). Practical selection of SVM parameters and noise estimation for SVM  
1139 regression. *Neural Networks*, 17(1), 113–126.
- 1140 Chua, L. H., & Wong, T. S. (2010). Improving event-based rainfall–runoff modeling using a combined  
1141 artificial neural network–kinematic wave approach. *Journal of Hydrology*, 390(1-2), 92-107.
- 1142 Corzo, G. A., Solomatine, D. P., Hidayat, de Wit, M., Werner, M., Uhlenbrook, S., and Price, R. K. (2009).  
1143 Combining semi-distributed process-based and data-driven models in flow simulation: a case study  
1144 of the Meuse river basin. *Hydrology and Earth System Sciences*, 13, 1619-1634,  
1145 <https://doi.org/10.5194/hess-13-1619-2009>.
- 1146 Dangeti, P. (2017). *Statistics for Machine Learning*. Packt Publishing Ltd.
- 1147 de Villiers, J., and Barnard, E. (1993). Backpropagation neural nets with one and two hidden layers. *IEEE*  
1148 *Transactions on Neural Networks*, 4(1), 136-141.
- 1149 DeBeer, C. M., Wheeler, H. S., Pomeroy, J. W., Barr, A. G., Baltzer, J. L., Johnstone, J. F., Turetsky, M. R.,  
1150 Stewart, R. E., Hayashi, M., van der Kamp, G., Marshall, S., Campbell, E., Marsh, P., Carey, S. K.,  
1151 Quinton, W. L., Li, Y., Razavi, S., Berg, A., McDonnell, J. J., Spence, C., Helgason, W. D., Ireson, A. M.,  
1152 Black, T. A., Davison, B., Howard, A., Thériault, J. M., Shook, K., and Pietroniro, A. (2020). Summary  
1153 and synthesis of Changing Cold Regions Network (CCRN) research in the interior of western Canada  
1154 – Part 2: Future change in cryosphere, vegetation, and hydrology. *Hydrology and Earth System*  
1155 *Sciences Discussions*, <https://doi.org/10.5194/hess-2020-491>.
- 1156 Dechter, R. (1986). *Learning while searching in constraint-satisfaction problems*. University of California,  
1157 Computer Science Department, Cognitive Systems Laboratory.
- 1158 Deist, T. M., Patti, A., Wang, Z., Krane, D., Sorenson, T., & Craft, D. (2019). Simulation-assisted machine  
1159 learning. *Bioinformatics*, 35(20), 4072-4080.
- 1160 Dengiz, B., Alabas-Uslu, C., & Dengiz, O. (2009). A tabu search algorithm for the training of neural  
1161 networks. *Journal of the Operational Research Society*, 60(2), 282-291.
- 1162 Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional  
1163 transformers for language understanding. arXiv preprint arXiv:1810.04805.
- 1164 Ducournau, A., & Fablet, R. (2016, December). Deep learning for ocean remote sensing: an application of  
1165 convolutional neural networks for super-resolution on satellite-derived SST data. In *2016 9th IAPR*  
1166 *Workshop on Pattern Recognition in Remote Sensing (PRRS)* (pp. 1-6). IEEE.
- 1167 Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, 14(2), 179-211.
- 1168 Eyring, V., Bony, S., Meehl, G. A., Senior, C. A., Stevens, B., Stouffer, R. J., & Taylor, K. E. (2016). Overview  
1169 of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and  
1170 organization. *Geoscientific Model Development*, 9(5), 1937-1958.
- 1171 Fama, E. F. (1970). Efficient capital markets: a review of theory and empirical work. *The Journal of Finance*,  
1172 25(2), 383-417.
- 1173 Fang, X., & Pomeroy, J. W. (2020). Diagnosis of future changes in hydrology for a Canadian Rockies  
1174 headwater basin. *Hydrology and Earth System Sciences*, 24(5), 2731-2754.

- 1175 Foresee, F. D., & Hagan, M. T. (1997, June). Gauss-Newton approximation to Bayesian learning. In  
1176 Proceedings of International Conference on Neural Networks (ICNN'97) (Vol. 3, pp. 1930-1935). IEEE.
- 1177 Frasconi, P., Gori, M., & Soda, G. (1992). Local feedback multilayered networks. *Neural Computation*, 4(1),  
1178 120-130.
- 1179 Friedman, J. H. (1991). Multivariate adaptive regression splines. *Annals of Statistics*, 19(1), 1–67.
- 1180 Gardner, M. W., & Dorling, S. R. (1998). Artificial neural networks (the multilayer perceptron)—a review  
1181 of applications in the atmospheric sciences. *Atmospheric Environment*, 32(14-15), 2627-2636.
- 1182 Gelaro, R., McCarty, W., Suárez, M. J., Todling, R., Molod, A., Takacs, L., ... & Wargan, K. (2017). The  
1183 modern-era retrospective analysis for research and applications, version 2 (MERRA-2). *Journal of*  
1184 *Climate*, 30(14), 5419-5454.
- 1185 Gharari, S., & Razavi, S. (2018). A review and synthesis of hysteresis in hydrology and hydrological  
1186 modeling: memory, path-dependency, or missing physics? *Journal of Hydrology*, 566, 500-519.
- 1187 Glorot, X., Bordes, A., & Bengio, Y. (2011, June). Deep sparse rectifier neural networks. In *Proceedings of*  
1188 *the 14<sup>th</sup> International Conference on Artificial Intelligence and Statistics* (pp. 315-323).
- 1189 Goodfellow, I., Bengio, Y., Courville, A., & Bengio, Y. (2016). *Deep learning* (Vol. 1). Cambridge: MIT Press.
- 1190 Guillaume, J. H., Jakeman, J. D., Marsili-Libelli, S., Asher, M., Brunner, P., Croke, B., ... & Stigter, J. D. (2019).  
1191 Introductory overview of identifiability analysis: a guide to evaluating whether you have the right  
1192 type of data for your modeling purpose. *Environmental Modelling & Software*, 119, 418-432.
- 1193 Gupta, H. V., Clark, M. P., Vrugt, J. A., Abramowitz, G., & Ye, M. (2012). Towards a comprehensive  
1194 assessment of model structural adequacy. *Water Resources Research*, 48(8).
- 1195 Hagan, M. T., & Menhaj, M. B. (1994). Training feedforward networks with the Marquardt algorithm. *IEEE*  
1196 *Transactions on Neural Networks*, 5(6), 989-993.
- 1197 Hendler, J. (2008). Avoiding another AI winter. *IEEE Intelligent Systems*, (2), 2-4.
- 1198 Hinton, G. E., McClelland, J., & Rumelhart, D. (1986). Distributed representations. In D. E. Rumelhart & J.  
1199 L. McClelland, editors, *Parallel Distributed Processing: Explorations in the Microstructure of*  
1200 *Cognition*, volume 1, pages 77–109. MIT Press, Cambridge.
- 1201 Hinton, G. E., Osindero, S., & Teh, Y. W. (2006). A fast learning algorithm for deep belief nets. *Neural*  
1202 *Computation*, 18(7), 1527-1554.
- 1203 Hinton, G. E., Srivastava, N., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. R. (2012). Improving neural  
1204 networks by preventing co-adaptation of feature detectors. arXiv preprint arXiv:1207.0580.
- 1205 Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735-1780.
- 1206 Hornberger, G. M., & Spear, R. C. (1981). Approach to the preliminary analysis of environmental systems.  
1207 *J. Environ. Mgmt.*, 12(1), 7-18.
- 1208 Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal  
1209 approximators. *Neural Networks*, 2(5), 359-366.

- 1210 Hosny, A., Parmar, C., Quackenbush, J., Schwartz, L. H., & Aerts, H. J. (2018). Artificial intelligence in  
1211 radiology. *Nature Reviews Cancer*, 18(8), 500-510.
- 1212 Hsu, K. L., Gao, X., Sorooshian, S., & Gupta, H. V. (1997). Precipitation estimation from remotely sensed  
1213 information using artificial neural networks. *Journal of Applied Meteorology*, 36(9), 1176-1190.
- 1214 Hsu, K. L., Gupta, H. V., & Sorooshian, S. (1995). Artificial neural network modeling of the rainfall-runoff  
1215 process. *Water Resources Research*, 31(10), 2517-2530.
- 1216 Hutson, M. (2018). Has artificial intelligence become alchemy? *Science*, 360(6388), 478. DOI:  
1217 10.1126/science.360.6388.478
- 1218 Johnson, V. M., & Rogers, L. L. (2000). Accuracy of neural network approximators in simulation-  
1219 optimization. *Journal of Water Resources Planning and Management*, 126(2), 48-56.
- 1220 Jordan, M. I. (1986). Attractor dynamics and parallelism in a connectionist sequential machine. In 8th  
1221 Annual Conference, Cognitive Science Society. MIT Press: Amherst, MA; 531–546.
- 1222 Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: trends, perspectives, and prospects. *Science*,  
1223 349(6245), 255-260.
- 1224 Kang, K. W., Park, C. Y., & Kim, J. H. (1993). Neural network and its application to rainfall-runoff forecasting.  
1225 *Korean Journal of Hydrosiences*, 4, 1-9.
- 1226 Karamouz, M., Razavi, S., & Araghinejad, S. (2008). Long-lead seasonal rainfall forecasting using time-delay  
1227 recurrent neural networks: a case study. *Hydrological Processes: An International Journal*, 22(2), 229-  
1228 241.
- 1229 Karpatne, A., Atluri, G., Faghmous, J. H., Steinbach, M., Banerjee, A., Ganguly, A., ... & Kumar, V. (2017).  
1230 Theory-guided data science: a new paradigm for scientific discovery from data. *IEEE Transactions on*  
1231 *Knowledge and Data Engineering*, 29(10), 2318-2331.
- 1232 Kennedy, M. C., & O’Hagan, A. (2000). Predicting the output from a complex computer code when fast  
1233 approximations are available. *Biometrika*, 87(1), 1–13.
- 1234 Kim, S. S. (1998). Time-delay recurrent neural network for temporal correlations and prediction.  
1235 *Neurocomputing*, 20(1-3), 253-263.
- 1236 Kirchner, J. W. (2006). Getting the right answers for the right reasons: linking measurements, analyses,  
1237 and models to advance the science of hydrology. *Water Resources Research*, 42(3).
- 1238 Klemeš, V. (1986a). Operational testing of hydrological simulation models. *Hydrological Sciences Journal*,  
1239 31(1), 13-24.
- 1240 Klemeš, V. (1986b). Dilettantism in hydrology: transition or destiny? *Water Resources Research*, 22(9S),  
1241 177S-188S.
- 1242 Kolakowski, M. (2018). How Algo Trading Is Worsening Stock Market Routs, Investopedia  
1243 (<https://www.investopedia.com/news/how-algo-trading-worsening-stock-market-routs/>)

- 1244 Krasnopolsky, V. M. (2007). Neural network emulations for complex multidimensional geophysical  
1245 mappings: applications of neural network techniques to atmospheric and oceanic satellite retrievals  
1246 and numerical modeling. *Reviews of Geophysics*, 45(3).
- 1247 Kratzert, F., Klotz, D., Brenner, C., Schulz, K., & Herrnegger, M. (2018). Rainfall–runoff modelling using long  
1248 short-term memory (LSTM) networks. *Hydrology and Earth System Sciences*, 22(11), 6005-6022.
- 1249 Kratzert, F., Klotz, D., Herrnegger, M., Sampson, A. K., Hochreiter, S., & Nearing, G. S. (2019). Toward  
1250 improved predictions in ungauged basins: Exploiting the power of machine learning. *Water  
1251 Resources Research*, 55(12), 11344-11354.
- 1252 Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2017). Imagenet classification with deep convolutional neural  
1253 networks. *Communications of the ACM*, 60(6), 84-90.
- 1254 Krogh, A., & Hertz, J. (1991). A simple weight decay can improve generalization. *Advances in Neural  
1255 Information Processing Systems*, 4, 950-957.
- 1256 Lawrence, S., Giles, C. L., Tsoi, A. C., & Back, A. D. (1997). Face recognition: a convolutional neural-network  
1257 approach. *IEEE Transactions on Neural Networks*, 8(1), 98-113.
- 1258 LeCun, Y. (2017). My take on Ali Rahimi's "Test of Time" award talk at NIPS, (accessed online in December  
1259 2020 at: [https://www2.isye.gatech.edu/~tzhao80/Yann\\_Response.pdf](https://www2.isye.gatech.edu/~tzhao80/Yann_Response.pdf))
- 1260 Lee, J., Kim, R., Koh, Y., & Kang, J. (2019). Global stock market prediction based on stock chart images  
1261 using deep Q-network. *IEEE Access*, 7, 167260-167277.
- 1262 Lindström, G., Johansson, B., Persson, M., Gardelin, M., & Bergström, S. (1997). Development and test of  
1263 the distributed HBV-96 hydrological model. *Journal of hydrology*, 201(1-4), 272-288.
- 1264 MacKay, D. J. (1992). A practical Bayesian framework for backpropagation networks. *Neural Computation*,  
1265 4(3), 448-472.
- 1266 Maier, H. R., & Dandy, G. C. (2000). Neural networks for the prediction and forecasting of water resources  
1267 variables: a review of modelling issues and applications. *Environmental Modelling & Software*, 15(1),  
1268 101-124.
- 1269 Maier, H. R., Jain, A., Dandy, G. C., & Sudheer, K. P. (2010). Methods used for the development of neural  
1270 networks for the prediction of water resource variables in river systems: current status and future  
1271 directions. *Environmental Modelling & Software*, 25(8), 891-909.
- 1272 Malkiel, B. G. (2003). The efficient market hypothesis and its critics. *Journal of Economic Perspectives*,  
1273 17(1), 59-82.
- 1274 McCann, D. W. (1992). A neural network short-term forecast of significant thunderstorms. *Weather and  
1275 Forecasting*, 7(3), 525-534.
- 1276 Milly, P. C. D., Betancourt, J., Falkenmark, M., Hirsch, R. M., Kundzewicz, Z. W., Lettenmaier, D. P., &  
1277 Stouffer, R. J. (2008). Stationarity is dead: whither water management? *Science*, 319, 573-574.
- 1278 Minns, A. W., & Hall, M. J. (1996). Artificial neural networks as rainfall-runoff models. *Hydrological  
1279 Sciences Journal*, 41(3), 399-417.

- 1280 Minsky, M., & Papert, S. A. (1969). *Perceptrons: An Introduction to Computational Geometry*. MIT Press.
- 1281 Nash, J. E., & Sutcliffe, J. V. (1970). River flow forecasting through conceptual models part I—A discussion  
1282 of principles. *Journal of Hydrology*, 10(3), 282-290.
- 1283 Navone, H. D., & Ceccatto, H. A. (1994). Predicting Indian monsoon rainfall: a neural network approach.  
1284 *Climate Dynamics*, 10(6-7), 305-312.
- 1285 Nearing, G. S., Kratzert, F., Sampson, A. K., Pelissier, C. S., Klotz, D., Frame, J. M., ... & Gupta, H. V. (2020).  
1286 What role does hydrological science play in the age of machine learning?. *Water Resources Research*,  
1287 e2020WR028091.
- 1288 Oreskes, N., Shrader-Frechette, K., & Belitz, K. (1994). Verification, validation, and confirmation of  
1289 numerical models in the earth sciences. *Science*, 263(5147), 641-646.
- 1290 Panchal, J. H., Fuge, M., Liu, Y., Missoum, S., & Tucker, C. (2019). Machine learning for engineering design.  
1291 *Journal of Mechanical Design*, 141(11).
- 1292 Pearlstein, S. (2018). The robots-vs.-robots trading that has hijacked the stock market, *The Washington*  
1293 *Post*, [https://www.washingtonpost.com/news/wonk/wp/2018/02/07/the-robots-v-robots-trading-  
1294 that-has-hijacked-the-stock-market/](https://www.washingtonpost.com/news/wonk/wp/2018/02/07/the-robots-v-robots-trading-that-has-hijacked-the-stock-market/)
- 1295 Prechelt, L. (1998). Early stopping-but when? In G. Montavon, G. B. Orr, & K.-R. Müller, editors, *Neural*  
1296 *Networks: Tricks of the Trade*, pages 55-69. Springer, Berlin, Heidelberg.
- 1297 Rahimi, A., and & Recht, B., (2017)., Back when we were kids, Test-of-time award presentation,  
1298 Conference on Neural Information Processing Systems, 2017, Vancouver, Canada (Video online  
1299 accessed in December, 2020, at from: <https://youtu.be/Qi1Yry33TQE>).
- 1300 Raina, R., Madhavan, A., & Ng, A. Y. (2009, June). Large-scale deep unsupervised learning using graphics  
1301 processors. In *Proceedings of the 26th Annual International Conference on Machine Learning* (pp.  
1302 873-880).
- 1303 Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: a deep learning  
1304 framework for solving forward and inverse problems involving nonlinear partial differential  
1305 equations. *Journal of Computational Physics*, 378, 686-707.
- 1306 Rajbhandari, S., Rasley, J., Ruwase, O., & He, Y. (2019). ZeRO: memory optimization towards training a  
1307 trillion parameter models. arXiv preprint arXiv:1910.02054.
- 1308 Rakitianskaia, A., & Engelbrecht, A. P. (2009). Training neural networks with PSO in dynamic environments.  
1309 In *2009 IEEE Congress on Evolutionary Computation* (pp. 667-673).
- 1310 Razavi, S., & Araghinejad, S. (2009). Reservoir inflow modeling using temporal neural networks with  
1311 forgetting factor approach. *Water Resources Management*, 23(1), 39-55.
- 1312 Razavi, S., & Tolson, B. A. (2011). A new formulation for feedforward neural networks. *IEEE Transactions*  
1313 *on Neural Networks*, 22(10), 1588-1598.
- 1314 Razavi, S., Elshorbagy, A., Wheeler, H., & Sauchyn, D. (2015). Toward understanding nonstationarity in  
1315 climate and hydrology through tree ring proxy records. *Water Resources Research*, 51(3), 1813-1830.

- 1316 Razavi, S., Jakeman, A., Saltelli, A., Prieur, C., looss, B., Borgonovo, E., ... & Tarantola, S. (2021). The fFuture  
1317 of sSensitivity aAnalysis: aAn eEssential dDiscipline for sSystems mModeling and pPolicy sSupport.  
1318 Environmental Modelling & Software, 104954.
- 1319 Razavi, S., & Karamouz, M. (2007). Adaptive neural networks for flood routing in river systems. Water  
1320 international, 32(3), 360-375.
- 1321 Razavi, S., Sheikholeslami, R., Gupta, H. V., & Haghnegahdar, A. (2019). VARS-TOOL: A toolbox for  
1322 comprehensive, efficient, and robust sensitivity and uncertainty analysis. Environmental Modelling  
1323 & Software, 112, 95-107.
- 1324 Razavi, S., Tolson, B. A., & Burn, D. H. (2012a). Review of surrogate modeling in water resources. Water  
1325 Resources Research, 48(7).
- 1326 Razavi, S., Tolson, B. A., & Burn, D. H. (2012b). Numerical assessment of metamodelling strategies in  
1327 computationally intensive optimization. Environmental Modelling & Software, 34, 67-86.
- 1328 Reed, R. (1993). Pruning algorithms—A survey. IEEE Transactions on Neural Networks, 4(5), 740–747.
- 1329 Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., & Carvalhais, N. (2019). Deep learning  
1330 and process understanding for data-driven Earth system science. Nature, 566(7743), 195-204.
- 1331 Rodriguez-Iturbe, I., Entekhabi, D., Lee, J. S., & Bras, R. L. (1991). Nonlinear dynamics of soil moisture at  
1332 climate scales: 2. Chaotic analysis. Water Resources Research, 27(8), 1907-1915.
- 1333 Rosenblatt, F. (1957). The perceptron, a perceiving and recognizing automaton (Project PARA), Cornell  
1334 Aeronautical Laboratory Report No. 85-460-1, Buffalo, New York.
- 1335 Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating  
1336 errors. Nature, 323(6088), 533-536.
- 1337 Scheffer, M., Carpenter, S., Foley, J. A., Folke, C., & Walker, B. (2001). Catastrophic shifts in ecosystems.  
1338 Nature, 413(6856), 591-596.
- 1339 Schmidhuber, J. (2015a). "Deep Learning". Scholarpedia. 10(11), 32832. Bibcode:2015SchpJ..1032832S.  
1340 doi:10.4249/scholarpedia.32832.
- 1341 Schmidhuber, J. (2015b). Deep learning in neural networks: an overview. Neural Networks, 61, 85-117.
- 1342 Shamseldin, A. Y., & O'Connor, K. M. (2001). A non-linear neural network technique for updating of river  
1343 flow forecasts. Hydrology and Earth System Sciences, 5, 577–598, <https://doi.org/10.5194/hess-5-577-2001>.  
1344
- 1345 Shen, C. (2018). A transdisciplinary review of deep learning research and its relevance for water resources  
1346 scientists. Water Resources Research, 54(11), 8558-8593.
- 1347 Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., ... & Lillicrap, T. (2018). A general  
1348 reinforcement learning algorithm that masters chess, shogi, and Go through self-play. Science,  
1349 362(6419), 1140-1144.

- 1350 Sorooshian, S., Hsu, K. L., Gao, X., Gupta, H. V., Imam, B., & Braithwaite, D. (2000). Evaluation of PERSIANN  
1351 system satellite-based estimates of tropical rainfall. *Bulletin of the American Meteorological Society*,  
1352 81(9), 2035-2046.
- 1353 Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: a simple way  
1354 to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1), 1929-  
1355 1958.
- 1356 Stewart, R., & Ermon, S. (2017, February). Label-free supervision of neural networks with physics and  
1357 domain knowledge. In 31<sup>st</sup> AAAI Conference on Artificial Intelligence.
- 1358 Stogryn, A. P., Butler, C. T., & Bartolac, T. J. (1994). Ocean surface wind retrievals from special sensor  
1359 microwave imager data with neural networks. *Journal of Geophysical Research: Oceans*, 99(C1), 981-  
1360 984.
- 1361 Tamura, S., & Tateishi, M. (1997). Capabilities of a four-layered feedforward neural network: four layers  
1362 versus three. *IEEE Transactions on Neural Networks*, 8(2), 251–255.
- 1363 Tang, G., Long, D., Behrangi, A., Wang, C., & Hong, Y. (2018). Exploring deep neural networks to retrieve  
1364 rain and snow in high latitudes using multisensor and reanalysis data. *Water Resources Research*,  
1365 54(10), 8253-8278.
- 1366 Teoh, E. J., Tan, K. C., & Xiang, C. (2006). Estimating the number of hidden neurons in a feedforward  
1367 network using the singular value decomposition. *IEEE Transactions on Neural Networks*, 17(6), 1623-  
1368 1629.
- 1369 Tickle, A. B., Andrews, R., Golea, M., & Diederich, J. (1998). The truth will come to light: directions and  
1370 challenges in extracting the knowledge embedded within trained artificial neural networks. *IEEE*  
1371 *Transactions on Neural Networks*, 9(6), 1057-1068.
- 1372 Tokar, A. S., & Markus, M. (2000). Precipitation-runoff modeling using artificial neural networks and  
1373 conceptual models. *Journal of Hydrologic Engineering*, 5(2), 156-161.
- 1374 Torresen, J. (2018). A review of future and ethical perspectives of robotics and AI. *Frontiers in Robotics*  
1375 *and AI*, 4, 75.
- 1376 Towell, G. G., & Shavlik, J. W. (1994). Knowledge-based artificial neural networks. *Artificial Intelligence*,  
1377 70(1-2), 119-165.
- 1378 Vapnik, V. (1998) *Statistical Learning Theory*, Wiley, New York.
- 1379 von Rueden, L., Mayer, S., Beckh, K., Georgiev, B., Giesselbach, S., Heese, R., ... & Walczak, M. (2019).  
1380 Informed machine learning—A taxonomy and survey of integrating knowledge into learning systems.  
1381 arXiv preprint arXiv:1903.12394.
- 1382 Waibel, A., Hanazawa, T., Hintin, G., Shikano, K., Lang, K. J. (1989). Phoneme recognition using time delay  
1383 neural networks. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 37(3), 328–339.
- 1384 Wexler, R. (2017). When a Computer Program Keeps You in Jail. *The New York Times*  
1385 ([https://www.nytimes.com/2017/06/13/opinion/how-computers-are-harming-criminal-](https://www.nytimes.com/2017/06/13/opinion/how-computers-are-harming-criminal-justice.html?auth=login-google)  
1386 [justice.html?auth=login-google](https://www.nytimes.com/2017/06/13/opinion/how-computers-are-harming-criminal-justice.html?auth=login-google))

- 1387 Wilby, R. L., Abrahart, R. J., & Dawson, C. W. (2003). Detection of conceptual model rainfall—runoff  
1388 processes inside an artificial neural network. *Hydrological Sciences Journal*, 48(2), 163-181.
- 1389 Xiang, C., Ding, S. Q., & Lee, T. H. (2005). Geometrical interpretation and architecture selection of MLP.  
1390 *IEEE Transactions on Neural Networks*, 16(1), 84-96.
- 1391 Xingjian, S. H. I., Chen, Z., Wang, H., Yeung, D. Y., Wong, W. K., & Woo, W. C. (2015). Convolutional LSTM  
1392 network: a machine learning approach for precipitation nowcasting. In *Advances in Neural*  
1393 *Information Processing Systems* (pp. 802-810).
- 1394 Xu, Y., Wong, K. W., & Leung, C. S. (2006). Generalized RLS approach to the training of neural networks.  
1395 *IEEE Transactions on Neural Networks*, 17(1), 19–34.
- 1396 Young, T., Hazarika, D., Poria, S., & Cambria, E. (2018). Recent trends in deep learning based natural  
1397 language processing. *IEEE Computational Intelligence Magazine*, 13(3), 55-75.
- 1398