

# Detection of forced change within combined climate fields using explainable neural networks

Jamin K. Rader<sup>1</sup>, Elizabeth A. Barnes<sup>1</sup>, Imme Ebert-Uphoff<sup>2,3</sup>, Chuck  
Anderson<sup>4</sup>

<sup>1</sup>Department of Atmospheric Science, Colorado State University, Fort Collins, CO, USA

<sup>2</sup>Cooperative Institute for Research in the Atmosphere, Colorado State University, Fort Collins, CO, USA

<sup>3</sup>Department of Electrical and Computer Engineering, Colorado State University, Fort Collins, CO, USA

<sup>4</sup>Department of Computer Science, Colorado State University, Fort Collins, CO, USA

## Key Points:

- Neural networks and their explainability tools can be harnessed to identify patterns of forced change within combined fields
- Combined fields of input allow for earlier detection of the emergence of a forced climate response
- Explainable AI techniques can be used to identify patterns that describe the emergence and evolution of forced climate change

---

Corresponding author: Jamin K. Rader, [jamin.rader@colostate.edu](mailto:jamin.rader@colostate.edu)

**Abstract**

Assessing forced climate change requires the extraction of the forced signal from the background of climate noise. Traditionally, tools for extracting forced climate change signals have focused on one atmospheric variable at a time, however, using multiple variables can reduce noise and allow for easier detection of the forced response. Following previous work, we train artificial neural networks to predict the year of single- and multi-variable maps from forced climate model simulations. To perform this task, the neural networks learn patterns that allow them to discriminate between maps from different years—that is, the neural networks learn the patterns of the forced signal amidst the shroud of internal variability and climate model disagreement. When presented with combined input fields (multiple seasons, variables, or both), the neural networks are able to detect the signal of forced change earlier than when given single fields alone by utilizing complex, nonlinear relationships between multiple variables and seasons. We use layer-wise relevance propagation, a neural network explainability tool, to identify the multivariate patterns learned by the neural networks that serve as reliable indicators of the forced response. These “indicator patterns” vary in time and between climate models, providing a template for investigating inter-model differences in the time evolution of the forced response. This work demonstrates how neural networks and their explainability tools can be harnessed to identify patterns of the forced signal within combined fields.

**Plain Language Summary**

Using machine learning tools called neural networks, we identify patterns of the changing climate within climate model data. Changes in the climate can be identified earlier when detecting patterns within maps of multiple variables and seasons than for single maps alone. By visualizing the patterns learned by the neural networks, we can identify which regions, variables, and seasons are most important for detecting climate change. These patterns offer insight into how climate change is represented in different climate models, and how the patterns of climate change will evolve over time.

**1 Introduction**

Changes in the climate system comprise the Earth system’s response to anthropogenic external forcings (e.g. greenhouse gas and aerosol emissions), natural external forcings (e.g. variations in the solar cycle, volcanic activity), internal variability (natural variations in the climate due to internal processes), and the interactions between. Distinguishing which features of climate change are the product of external forcings, rather than a byproduct of internal variability, is critical for mitigation and adaptation science (Field

50 et al., 2014; Maher et al., 2021; Mankin et al., 2020; Sanderson et al., 2018). To iden-  
51 tify the forced response to external forcings, changes in the climate are often simplified  
52 into “signal” and “noise” components (e.g., Hawkins & Sutton, 2009; Mahony & Can-  
53 non, 2018; Scaife & Smith, 2018). The signal of climate change captures all anthropogenic  
54 and natural external forcings, which we refer to as the forced signal or forced response  
55 in this study. Climate noise, a combination of internal variability (natural variations in  
56 the climate due to internal processes) and climate model disagreement in the magnitude  
57 of the response, often acts to obscure the forced signal (Santer et al., 2011).

58 Innovative methods are required to determine which behaviors of the climate are  
59 the result of the forced signal and which are the result of climate noise. Decades of re-  
60 search have provided a diverse toolkit for this task (North & Stevens, 1998) which in-  
61 cludes linear regression (e.g., Mudelsee, 2019; Santer et al., 1996; Sippel et al., 2020; Solow,  
62 1987), empirical orthogonal functions and linear discriminant analysis (e.g., Santer et  
63 al., 2019; Schneider & Held, 2001; Wills et al., 2018, 2020), and linear inverse models (e.g.,  
64 Solomon & Newman, 2012), to name a few. Recently, neural networks have also entered  
65 the fold. Neural networks are machine learning algorithms that are able to detect com-  
66 plex, nonlinear relationships between input and output data (Abiodun et al., 2018). Be-  
67 cause neural networks are able to detect highly complex relationships, they are useful  
68 for many high dimensional problems and have become prevalent in several atmospheric  
69 science research fields, such as weather forecasting (e.g., Lagerquist et al., 2019; Lee et  
70 al., 2021; Weyn et al., 2020), climate model parameterizations (e.g., Brenowitz & Brether-  
71 ton, 2018; Gettelman et al., 2021; Silva et al., 2021), and, most relevant to the focus of  
72 this study, detection of a forced climate response (e.g., Barnes et al., 2019, 2020; Labe  
73 & Barnes, 2021; Madakumbura et al., 2021). To detect patterns of forced change, Barnes  
74 et al. (2020) trained a neural network to predict the year label of maps of annual-mean  
75 temperature (or precipitation) from climate model simulations for forced historical and  
76 future scenarios. Given that the internal variability in any given year differs between the  
77 various climate models, the neural network had to learn patterns of the forced climate  
78 response. Using neural network explainability methods, they then visualized the regions  
79 that were most reliable indicators for identifying change across the CMIP5 models. Barnes  
80 et al. (2020) demonstrated that neural networks, and their explainability methods, are  
81 powerful tools for extracting forced patterns from climate data. Since then, neural net-  
82 works have been used to explore the sensitivity of regional temperature signals to aerosols  
83 and greenhouse gases using single-forcing large ensembles, and to detect the signal of ex-  
84 treme precipitation in observational datasets (Labe & Barnes, 2021; Madakumbura et  
85 al., 2021).

86            Though many climate signal detection studies focus on single variables, such as annual-  
87 mean temperature or a single season of precipitation (Gaetani et al., 2020; Li et al., 2017;  
88 Santer et al., 1996, 2019), there are benefits to studying climate change through a mul-  
89 tivariate lens (Bindoff et al., 2013; Bonfils et al., 2020; Mahony & Cannon, 2018). Many  
90 variables in our atmosphere are closely interconnected so, when the variables are intel-  
91 ligently selected, signals of change within multiple variables may be detected earlier than  
92 in single variables alone. For example, departure from natural variability can be seen decades  
93 earlier in bivariate maps of summertime temperature and precipitation than in either  
94 variable alone (Mahony & Cannon, 2018). Similarly, Fischer and Knutti (2012) found  
95 that climate model biases in the signal of relative humidity and temperature are neg-  
96 atively correlated such that climate model simulations of their combined quantity, heat  
97 stress, have considerably less spread. Combined variables have also been used to iden-  
98 tify the impacts of anthropogenic forcings on climate in observational datasets by iden-  
99 tifying the multivariate patterns that enhance the signal of change relative to the un-  
100 derlying noise (e.g., Barnett et al., 2008; Marvel & Bonfils, 2013). Understanding how  
101 the patterns of the forced response take shape through multiple atmospheric variables  
102 also allows for a deeper understanding of the physics at play, as in Bonfils et al. (2020).  
103 They explored the evolution of the climate fingerprint by analyzing the leading combined  
104 empirical orthogonal functions of temperature, precipitation, and climate moisture in-  
105 dex. This multivariate approach illuminated two cross-variable patterns of change: in-  
106 tensification of wet-dry patterns and meridional shifts in the ITCZ associated with in-  
107 terhemispheric temperature contrasts. Neither pattern can be fully explained by a sin-  
108 gle variable which highlights the utility of combining variables when identifying patterns  
109 of the forced response.

110            Providing a method for both nonlinear and multi-variable analysis of the forced re-  
111 sponse, this study extends the neural-network approach of Barnes et al. (2020) to com-  
112 bined fields of input. Combined fields could mean the same variable for different tem-  
113 poral segments (e.g. seasons), or different geophysical variables, both of which are ex-  
114 plored here. While this study is consistent with the majority of the methodology con-  
115 tained in Barnes et al. (2020) there are some key differences, including the application  
116 to combined fields, the spatial domains (Section 2.1), how the data is standardized and  
117 the hyperparameters are chosen (Section 3.2), and the addition of clustering to better  
118 understand the multiple pathways of the neural-network decision-making process (Sec-  
119 tion 5). Section 2 outlines the climate models and observations analyzed in this study.  
120 Section 3 introduces the neural network design, the explainability technique (layer-wise  
121 relevance propagation; LRP), and their applications to detection of the forced climate  
122 response. We then apply these methods to global temperature and precipitation over land

123 in Section 4. Here we investigate the benefits of combining variables and compare the  
124 results of the neural network with the classical approach of calculating signal-to-noise  
125 ratios. In Section 5, we explore the patterns of the forced response for extreme precip-  
126 itation over the Americas and investigate the applications of LRP to studying the evo-  
127 lution of nonlinear climate patterns across multiple climate models. Finally, Section 6  
128 highlights the advantages and disadvantages of using neural networks and LRP for forced  
129 response detection and scientific exploration.

## 130 **2 Data**

### 131 **2.1 CMIP6 Climate Models**

132 We use climate model output from the sixth phase of the Coupled Model Intercom-  
133 parison Project (CMIP6; Eyring et al., 2016). Specifically we focus on monthly-, seasonal-  
134 , and annual-mean fields of 2-meter air temperature ( $K$ ), precipitation rate ( $kg\ m^{-2}\ s^{-1}$ ),  
135 and precipitation rate from very wet days ( $kg\ m^{-2}\ s^{-1}$ ), hereafter referred to as temper-  
136 ature, precipitation, and extreme precipitation, respectively. We use the meteorological  
137 seasons of December-January-February (DJF), March-April-May (MAM), June-July-August  
138 (JJA), and September-October-November (SON) for calculating seasonal-mean fields.  
139 Defining seasons in this way allows for the earliest detection of forced change (see Fig-  
140 ure S1 for more details).

141 Very wet days are defined as days that exceed the 95th percentile of all days with  
142 precipitation over a pre-defined baseline period (Donat et al., 2016). This is a popular  
143 index for measuring changes in extreme precipitation (Cui et al., 2019; Kim et al., 2020)  
144 and is used as an indicator of climate change in the U.S. Global Climate Research Pro-  
145 gram (USGCRP, 2018). We define the baseline as the 40 years from 1980 to 2019, a pe-  
146 riod for which daily precipitation data exists in both the climate models and the obser-  
147 vations. To remove the instances in which climate models simulate sub-trace daily pre-  
148 cipitation totals, we only include days that simulated at least 1 mm of precipitation when  
149 calculating the 95th percentile of all days with precipitation (Dai et al., 2007).

150 One ensemble member is selected for each of the 37 CMIP6 climate models ana-  
151 lyzed so each climate model is only represented once in the training and testing data.  
152 Since daily output is required to calculate very wet days, we are limited to 32 models  
153 for extreme precipitation (Figure S2). We analyze the climate model data from 1920 to  
154 2098 under historical forcing (1920–2014) and the SSP585 scenario (2015–2098). SSP585  
155 represents the highest development pathway within CMIP6 scenarios (O’Neill et al., 2016),

156 combining shared socioeconomic pathway 5 (SSP5) and representative concentration path-  
157 way 8.5 (RCP8.5).

158 Our neural network methodology requires that all climate model fields have the same  
159 shape. To accommodate this we regrid the climate model fields from their native res-  
160 olutions using the second-order conservative remapping method in the Climate Data Op-  
161 erators package from MPI (Schulzweida, 2019). For temperature and precipitation, the  
162 data is regridded to 4 degrees latitude by 4 degrees longitude. For extreme precipitation,  
163 the data is regridded to 1.5 degrees latitude by 1.5 degrees longitude. This regridding  
164 step reduces the spatial resolution of the data for most climate models.

165 Two spatial domains are considered in the results of this paper. For temperature  
166 and precipitation, the neural networks are trained on all land north of 60°S. Here, we  
167 choose to focus on land grid points because that is where humanity lives and will acutely  
168 feel the impacts of changing surface temperatures and precipitation. We also exclude Antarc-  
169 tica where climate models and reanalyses struggle to accurately simulate temperature  
170 and precipitation. Each map of temperature and precipitation has 948 unique data points.  
171 For extreme precipitation, the neural networks are trained on North and South Amer-  
172 ica (land grid points bounded by 90°N, 55°S, 170°W, and 25°W). Here, we choose to nar-  
173 row the regional scope to show that neural networks are powerful tools for identifying  
174 the forced response even when the spatial domain, and thus the available data, is lim-  
175 ited. Each map of extreme precipitation has 2314 unique data points.

## 176 **2.2 Observations**

177 While this work largely focuses on the results of the neural network fed climate model  
178 data as input, we show that these methods can be extended to observational data as well.  
179 For temperature, we use the Berkeley Earth Surface Temperature (BEST) dataset (Rohde  
180 & Hausfather, 2020). This dataset provides both a temperature climatology and the anoma-  
181 lies at monthly resolution from 1850 to the present. We added the anomalies to the cli-  
182 matology to reconstruct the absolute temperature ( $K$ ) at each grid point for all months  
183 between 1920 and 2019. Monthly observational precipitation fields are obtained from the  
184 NOAA Global Precipitation Climatology Project (GPCP), version 2.3, for 1979 to the  
185 present (Adler et al., 2018). Since daily GPCP precipitation observations are only avail-  
186 able back to October 1996, we use the European Centre for Medium-Range Weather Fore-  
187 casts' ERA5 global reanalysis (Hersbach et al., 2020) at 6-hour resolution to construct  
188 observational monthly mean extreme precipitation fields from 1980 to the present. All  
189 observations are regridded in the same way as the climate model data for each respec-  
190 tive variable.

### 3 Forced Change Detection Framework

#### 3.1 Neural Network Design

To identify indicator patterns of the forced response for combined fields we first develop artificial neural networks that, given maps of CMIP6 climate model output from every simulated year from 1920 to 2098, are tasked to predict the year that is being simulated. The results for neural networks trained on ten different input vectors are explored in the following two sections. The input vectors include annual-, seasonal-, and monthly-mean data for temperature, precipitation, and temperature and precipitation combined, as well as seasonal-mean maps for extreme precipitation over the Americas. We use this diverse selection of input vectors to compare neural network performance and indicator patterns for single-field and combined-field inputs.

The neural network architecture is illustrated in Figure 1. Each unit of the input layer vector to the neural network corresponds to a different grid point in the input fields. For example, a neural network that uses seasonal-mean maps of temperature and precipitation as input (two variables and four seasons for a total of eight maps, 948 grid points per map) would have an input vector with 7,584 units. In all cases, this input layer is followed by two fully connected hidden layers with ten nodes each. More complicated hidden layer architectures were explored, such as two layers with 20 nodes each and three layers with 10 nodes each, but the benefits were negligible. The hidden layers are followed by an output layer that consists of 22 classes, one corresponding to each decade midpoint between 1905 and 2115 (e.g. 1905, 1915, 1925, . . . , 2115). A softmax function is applied to the outputs to convert them to units of likelihood, where the sum of the output vector is one. More details on the neural network design can be found in the supplementary materials.

The neural network is tasked with “predicting the year” rather than “predicting the decade” as the output layer may suggest. To translate between decade midpoints and individual year labels, we use fuzzy encoding (Zadeh, 1965) such that each year can be mapped to one or more neighboring classes with varying degrees of membership (encoded as likelihood). This is different than traditional methods that would map each year to a single decade midpoint. In the traditional case, 2040 and 2049 would be considered to be members of the same class since they are in the same decade, and information would be lost as there is no way to distinguish whether the samples come from the beginning or the end of the decade. Using fuzzy encoding, this information of where a sample lies in each decade is retained. We use a triangular membership function (Zadeh, 1965) with a width equal to one decade such that each year has partial membership in one or two

226 neighboring decade classes, and the total membership sums to one. Following this method,  
 227 any year directly on a decade midpoint has membership in that class only while years  
 228 that fall between decade midpoints have membership in the two neighboring classes. The  
 229 year 1925, for example, is mapped to a likelihood of one for the class 1925 and a like-  
 230 lihood of zero in all other classes. The year 2078 is mapped to a likelihood of 0.7 for the  
 231 2075 class and a likelihood of 0.3 for the 2085 class. Note that decoding class likelihoods  
 232 back to their year is simply the decade-weighted sum of the likelihood:  $0.7 \times 2075 + 0.3$   
 233  $\times 2085 = 2078$ . A visualization of the encoding/decoding process can be found in Fig.  
 234 2 of Barnes et al. (2020).

### 235 **3.2 Neural Network Training**

236 For each input vector we train 100 neural networks that differ only in which cli-  
 237 mate models are randomly split into the training and testing sets. Partitioning so that  
 238 each climate model’s samples are all part of either the training set or the testing set avoids  
 239 issues with autocorrelation (i.e. near-identical data appearing in both the training and  
 240 testing sets). One hundred neural networks provide a range of results across multiple com-  
 241 binations of training and testing simulations and offer confidence that the results are con-  
 242 sistent across CMIP6 climate models and do not overfit to any one training set. Each  
 243 neural network is trained over the entire 1920-2098 period on 80% of the climate model  
 244 simulations, and then tested on the remaining 20%. This leads to a training set of 30  
 245 simulations and a testing set of 7 simulations for temperature and precipitation fields,  
 246 and a training set of 26 simulations and a testing set of 6 simulations for extreme pre-  
 247 cipitation fields. We train the neural networks using the binary cross-entropy loss (see  
 248 Barnes et al., 2020) between the predicted class likelihoods and the correct class mem-  
 249 bership weights, such that the loss function is minimized when the two are equal. Prop-  
 250 erties of the neural network training process, such as the learning rate and activation func-  
 251 tions, can be found in the supplementary materials.

252 The neural networks have several hidden nodes which enable them to learn com-  
 253 plicated relationships between the input and output data. However, with limited train-  
 254 ing data, many of these learned relationships will capture patterns of the noise in the  
 255 training dataset which can lead to overfitting (Srivastava et al., 2014). To reduce over-  
 256 fitting, we apply ridge regularization ( $L_2$  regularization, see Barnes et al., 2020) to the  
 257 weights of the first hidden layer. Ridge regularization adds a penalty (called the ridge  
 258 penalty) to the square of the weights so the solution is penalized for having large weights.  
 259 Through training, this acts to shrink the largest weights, thus spreading the weight out  
 260 more evenly across multiple grid points. In our application this results in a more even

261 distribution of weight across regions with strong spatial correlation and improves the per-  
 262 formance of the neural networks when given data they were not trained on, namely those  
 263 models in the testing set. Unlike classical approaches which tune the neural network to  
 264 reduce the mean squared error (MSE) between the predicted and truth outputs in the  
 265 testing set (in our case this would be the MSE between the truth and predicted years),  
 266 we select the ridge penalty that minimizes the time of emergence of the forced climate  
 267 signal (see Section 3.3). Using time of emergence, rather than MSE, to identify the ap-  
 268 propriate ridge penalty ensures that we are encouraging the neural networks to learn the  
 269 patterns of the forced response across all decades. When a small ridge penalty is used,  
 270 the neural networks are able to predict the year at the end of the 21st century almost  
 271 perfectly, at the expense of the predictive skill in earlier decades. This results in a later  
 272 calculation of time of emergence for the testing set. Increasing the ridge penalty allows  
 273 the neural networks to detect the climate change signal earlier (Figure S3). The ridge  
 274 penalty used for each input vector can be found in the supplementary materials. We use  
 275 the same ridge penalty for all 100 neural networks trained on each input vector.

276 All input fields (for climate models and observations) are standardized to assist with  
 277 the training and overall performance of the neural network. We subtracted the 1980–2019  
 278 mean at each grid point of the input fields for each climate model independently. This  
 279 recasts each input field to measure the change relative to the 1980–2019 mean, rather  
 280 than the raw magnitudes, which is also appropriate for identifying indicator patterns of  
 281 forced change. Since values for precipitation change are often on the order of  $10^{-6}$ , while  
 282 the values for temperature change are on the order of  $10^0$ , we normalized the data so the  
 283 inputs to the neural network all have a similar magnitude. To do this, the data from 1980–2019  
 284 at each grid point for each climate model are detrended using ordinary least squares lin-  
 285 ear regression. We then take the multi-model mean of the standard deviation of the de-  
 286 trended 1980–2019 data for each grid point. The input fields are then divided by this  
 287 new field of standard deviations so the inputs are of the same magnitude and fall in a  
 288 reasonable range for training the neural networks. Since all our observational datasets  
 289 include the years 1980 to 2019, we standardize the observations as if they were additional  
 290 climate models: raw observations are subtracted by their own 1980–2019 mean, and di-  
 291 vided by the same multi-model standard deviations that were used to standardize the  
 292 CMIP6 data.

### 293 **3.3 Time of Emergence Calculation**

294 The time of emergence of the forced climate response (hereafter, simply “TOE”)  
 295 is the time in which the forced response signal is distinguishable from the background

296 climate by the neural network. Specifically, we define the TOE as the year when the neu-  
 297 ral network is able to distinguish that year’s map from any map over a histor-  
 298 ical baseline period. In this work we define this baseline period as 1920–1959 and, under this def-  
 299 inition, the earliest possible TOE estimate is 1960. The TOE is estimated for each cli-  
 300 mate model simulation independently and a schematic of how the TOE is estimated is  
 301 presented in Figure 2. First, we calculate the maximum of the neural network-predicted  
 302 years over 1920–1959 for each model, which is referred to as the baseline maximum. We  
 303 then identify the TOE as the earliest year in which a map, and all subsequent maps, per-  
 304 manently exceed the baseline maximum. In Figure 2, sample model 1 has a baseline max-  
 305 imum of 1966 and permanently exceeds this prediction threshold in 2028. Sample model  
 306 2 has a baseline maximum of 1981 and permanently exceeds this threshold in 1989. Thus,  
 307 the TOE for sample models 1 and 2 are estimated as 2028 and 1989, respectively. In the  
 308 following sections we present the TOE for the testing set, however TOE estimates are  
 309 similar for both the training and testing sets.

### 310 **3.4 Layer-wise Relevance Propagation**

311 To visualize the patterns learned by the neural network we apply layer-wise rele-  
 312 vance propagation (LRP) which highlights the regions that were most relevant in the neu-  
 313 ral network’s decision-making process (Bach et al., 2015; Montavon et al., 2019). Toms  
 314 et al. (2020) discusses in detail how LRP can be used for neural network explainability  
 315 in the geosciences, though the most relevant details of LRP are described here.

316 LRP is a neural network explainability method that traces how information flows  
 317 through the pathways of a trained neural network. The values in a single-sample input  
 318 vector (in our case, a single year) are passed forward through the neural network. Us-  
 319 ing the same weights and activations used in the forward pass, LRP then propagates a  
 320 single-valued output back through the neural network to infer the extent to which each  
 321 of the values in the input layer contribute to the output (see Fig. 2 in Bach et al., 2015).  
 322 We refer to this quantity as relevance. Through this backpropagation process the out-  
 323 put value is conserved such that the sum of all relevance is equal to the output. At first  
 324 order, relevance can be likened to the product of the regression weights and input map  
 325 in a linear model. This quantity is natively unitless, but we convert it to a fraction by  
 326 dividing by the output value. This way, we can consider the relevance of a single pixel  
 327 in terms of its fractional contribution to the predicted class. Since LRP propagates only  
 328 a single output value at a time, we propagate relevance only for the decade class with  
 329 the highest likelihood. While the relevance maps detected by these networks evolve from

330 year to year, this evolution is slow so we find visualizing the highest likelihood decade  
 331 is sufficient.

332 There are several LRP decomposition rules which provide different methods of vi-  
 333 sualizing neural networks (Lapuschkin, 2019; Mamalakis et al., 2021). In our applica-  
 334 tions we use the  $\alpha\beta$ -rule which propagates positive relevance (regions that act to increase  
 335 the class likelihood) and negative relevance (regions that act to decrease the class like-  
 336 lihood) separately. Using the parameters  $\alpha = 1$  and  $\beta = 0$  we choose to only propa-  
 337 gate positive relevance, thus highlighting the regions that added to the likelihood of the  
 338 selected decade class.

### 339 **3.5 Signal-to-Noise Ratio Calculation**

340 In Section 4, we compare the LRP relevance maps to maps of signal-to-noise ratio  
 341 (S/N ratio), a more conventional method for identifying indicator patterns of the forced  
 342 response. S/N ratio consists of three distinct components: the forced signal, which is di-  
 343 vided by the sum of noise due to internal variability, and noise due to climate model dis-  
 344 agreement. A higher S/N ratio indicates that the signal of the forced response within  
 345 the climate models is very large relative to the underlying noise. We evaluate the S/N  
 346 ratio for each grid point separately, following the methodology in Hawkins and Sutton  
 347 (2012). First, we smooth the data from 1920 to 2098 for each climate model using a fourth-  
 348 order polynomial fit. The signal is defined as the difference between 2090 and 1920 in  
 349 the smoothed data, while internal variability is defined as the standard deviation of the  
 350 residuals from the smoothed data, and climate model disagreement is defined as the stan-  
 351 dard deviation of the signals calculated for all the climate models. S/N ratio is calcu-  
 352 lated by dividing the climate signal by the 90% confidence interval in the noise: inter-  
 353 nal variability and climate model disagreement. S/N ratio, and its components, can be  
 354 seen in Figure S7.

## 355 **4 Global Precipitation and Temperature**

### 356 **4.1 Time of Emergence**

357 Across all input vectors of temperature and precipitation, the neural networks are  
 358 able to learn patterns of the forced response. In the early-to-mid 20th century the forced  
 359 signal is small and undetectable by the neural networks amidst the noise of internal vari-  
 360 ability and model disagreement, which leads to poor predictive skill (Figure 3). How-  
 361 ever, as the signal increases in magnitude into the late-20th and 21st centuries, the neu-  
 362 ral networks are able to detect the patterns of the forced response and distinguish be-

363 tween maps in different years. These patterns of the forced response detected by the neu-  
364 ral networks are generalizable across CMIP6 models, and as a result the neural network  
365 has predictive skill for seen data (the training set) as well as unseen data (the testing  
366 set). These behaviors are shown in Figure 3 which presents the predicted years from one  
367 trained neural network for each combination of global precipitation and temperature in-  
368 put fields. Across all input vectors, a similar story of the forced signal unfolds. Prior to  
369 the TOE, the neural network is unable to identify patterns that allow it to accurately  
370 predict the year. As a result, the neural network is equally confident (or unconfident)  
371 about which year, between 1920 and the TOE, each input came from, so it predicts years  
372 right around the middle of the 20th century. After the TOE, the predicted years tend  
373 to follow a 1:1 line with the truth years, indicating that the neural network has identi-  
374 fied reliable indicators of change for this period.

375 Given that the predictive skill is similar for the training and testing sets we are con-  
376 fident that these neural networks, and thus their learned patterns, are generalizable across  
377 CMIP6 models. Although the neural networks are trained on climate model simulations,  
378 their learned patterns can be used to predict the year for observational data as well. When  
379 observations are used as input, the predicted years increase with time, just as they do  
380 for climate model input (Figure 3). This means that the indicators of change derived by  
381 the neural networks trained on climate models simulations are largely consistent with  
382 the real world. Pearson correlations ( $r$ ) of the actual years with the years predicted by  
383 each neural network are shown in Figure 4. All correlations are positive, indicating that  
384 the years predicted by the neural networks increase with time. These correlations are  
385 strongest for temperature and combined observations ( $r \approx 0.9$ ), but still quite high for  
386 precipitation ( $r \approx 0.8$ ). Correlations of actual years with predicted years are slightly higher  
387 for the combined temperature and precipitation observations than for temperature ob-  
388 servations alone (Figure S4), suggesting that the multivariate indicator patterns derived  
389 from climate model data are useful for understanding trends in the present-day climate.  
390 Across all variables, the highest observational correlations are found by the neural net-  
391 works trained on seasonal-mean data. The correlation of actual years with predicted years  
392 for precipitation observations are sensitive to the dataset of choice, which is expanded  
393 on in Section S3 and Figures S4 and S5.

394 The average TOEs, calculated from the climate models in the testing sets of all 100  
395 trained neural networks for each input field (Figure 5), reveal that the forced response  
396 can be detected earlier in maps of temperature than in maps of precipitation (Figure 5a-  
397 c). When presented with combined fields the neural networks are, in many cases, able  
398 to detect the forced signal even earlier than when given single fields alone (Figure 5b,f).

399 The TOE is generally earlier for the neural networks trained on seasonal-mean data than  
 400 for the neural networks trained on annual-mean data (Figure 5d-f). This is most notable  
 401 for precipitation fields, likely because there are large seasonal precipitation responses muted  
 402 by taking the annual mean (Tabari & Willems, 2018; Zappa et al., 2015). The TOEs are  
 403 earlier for temperature and precipitation combined than temperature alone when using  
 404 seasonal-mean maps (Figure 5b), but are approximately equal when using annual-mean  
 405 or monthly-mean maps (Figure 5a,c), which suggests that precipitation only improves  
 406 upon the detectability of the forced temperature signal when seasonal-mean fields are  
 407 used. While annual-mean precipitation may mute seasonal precipitation signals, monthly-  
 408 mean precipitation is noisy. In this case, seasonal means emerge as the appropriate tem-  
 409 poral segments for detecting precipitation change, underlining the importance for the  
 410 intentional and intelligent selection of neural network inputs.

411 The neural networks identify the earliest TOEs when trained on seasonal-mean tem-  
 412 perature and precipitation combined (Figure 5b,f). The TOE results for all 100 seasonal-  
 413 mean neural networks are summarized in the box plots in Figure S6. While the improve-  
 414 ment in forced response detection is small when precipitation is combined with temper-  
 415 ature, it is still notable given that the forced signal of temperature is much clearer than  
 416 the forced signal of precipitation. We use these variables as an initial example for em-  
 417 ploying this neural network methodology. We anticipate that more robust results might  
 418 be found for combinations of variables that have more distinct combined signals.

## 419 **4.2 Indicator Patterns for Combined Variables**

420 Having shown that the neural networks are able to predict the year given seasonal  
 421 means of temperature and precipitation (Figures 3, 5), we now identify and explore the  
 422 spatial indicator patterns used by the neural networks to make correct predictions. By  
 423 understanding the neural networks' decision-making process, we can identify which re-  
 424 gions act as combined (multi-seasonal and multi-variable) indicators of forced change amidst  
 425 a background of internal variability and climate model disagreement. To identify these  
 426 indicator patterns, we apply LRP to all climate model samples in the training and test-  
 427 ing sets from the year 2090 for the seasonal-mean combined neural networks. Averag-  
 428 ing the LRP results for each season and variable, we highlight the regions that have the  
 429 highest mean relevance across the 37 CMIP6 climate models and 100 trained neural net-  
 430 works. The relevance maps for temperature (precipitation) are shown in Figure 6a-d (7a-  
 431 d).

432 LRP identifies temperature over North Africa and Central Asia in JJA (Figure 6c)  
 433 and the Andes and Central Africa in SON (Figure 6d) as the most relevant regions for

434 predicting the year. For precipitation, the regions of highest relevance can be found in  
 435 Canada and Russia in DJF and SON (Figure 7a,d) and in Central Africa and India in  
 436 JJA and SON (Figure 7c,d). That is to say that these are the regional patterns iden-  
 437 tified by the neural networks that indicate the presence of forced change across the CMIP6  
 438 climate models. The scale of the color bars are different between Figures 6 and 7, such  
 439 that the darkest regions in the temperature maps are approximately one order of mag-  
 440 nitude more relevant than the darkest regions in the precipitation maps. Hence, the neu-  
 441 ral network is relying more heavily on the temperature inputs than the precipitation in-  
 442 puts in order to accurately predict the year. This is not surprising because the forced  
 443 signal of temperature is clearer than the forced signal of precipitation (Fig. SPM.7 in  
 444 Field et al., 2014). Even so, including seasonal precipitation allows the neural networks  
 445 to detect forced change earlier within combined fields than in temperature fields alone  
 446 (Figure 5b). The improvement in neural network performance provided by precipitation  
 447 (alongside temperature) is particularly noteworthy given that the S/N ratio for temper-  
 448 ature is larger than the S/N ratio for precipitation in all seasons and regions (Figures  
 449 6e-h, 7e-h, discussed further in this section). In other words, the forced temperature sig-  
 450 nal is always more pronounced than the forced precipitation signal, but the precipita-  
 451 tion signal is still useful for detecting forced change.

452 LRP is designed to highlight the regions that were most relevant for predicting the  
 453 correct class (in our case, the correct decade class). These LRP indicator patterns for  
 454 2090 are not the time-mean patterns of the forced response, they are the patterns used  
 455 by the neural network to distinguish the end of the 21st century from all other decades.  
 456 This is distinctly different from S/N ratio which identifies the regions where the forced  
 457 change from 1920 to 2090 is largest relative to internal variability and climate model spread.  
 458 Maps of S/N ratio for temperature are shown in Figure 6e-h, and the corresponding maps  
 459 for precipitation are shown in Figure 7e-h, where a higher S/N ratio (darker green) in-  
 460 dicates a clearer forced signal. For the most part, the indicator patterns identified by LRP  
 461 correspond with the regions with the highest S/N ratios. Calculating the Spearman's rank  
 462 correlation ( $\rho$ ) between each map of relevance and S/N ratio, we find that there is gen-  
 463 erally a strong positive correlation ( $0.71 \leq \rho \leq 0.77$ ) between the LRP indicator pat-  
 464 terns and the S/N ratios for temperature, and a moderate positive correlation ( $0.30 \leq$   
 465  $\rho \leq 0.56$ ) for precipitation. The exact correlation coefficients between each map are dis-  
 466 played in the subtitles for Figures 6e-h and 7e-h.

467 Given that precipitation only contributes a small amount of relevance compared  
 468 to temperature, it is perhaps unsurprising that there are several regions where the S/N  
 469 ratio for precipitation is high, but the relevance is low (e.g. Alaska in JJA, Figure 7c,g

470 or South Africa in SON, Figure 7d,h). Most likely, the forced signal of temperature is  
 471 clear enough that these regions do not add to the predictive skill of the neural networks.  
 472 Regions also exist where the S/N ratio for temperature is high despite low relevance (e.g.  
 473 North Africa in DJF, Figure 6a,e), although these are more rare, as hinted by the strong  
 474 correlation between the temperature maps of S/N ratio and relevance. In contrast, there  
 475 are fewer regions with high relevance despite low S/N ratios, but they do occur (e.g. SON  
 476 temperatures in northern South America, Figure 6d,h). These regions confirm that the  
 477 indicator patterns identified by LRP capture more than the local S/N ratio. They also  
 478 capture which regions improve detection of forced change only when coupled with sig-  
 479 nals in other regions, variables, and seasons. In the next section, we discuss further ap-  
 480 plications of neural network-derived indicator patterns and task the network with the  
 481 much harder problem of identifying changes in extreme precipitation over the Americas.

## 482 **5 Extreme Precipitation over the Americas**

483 We now task the neural networks to predict the year given combinations of sea-  
 484 sons for a single variable: extreme precipitation over the Americas. We choose to shift  
 485 our focus for a few reasons. First, we wish to demonstrate that this neural network ap-  
 486 proach can be extended to variables that have considerable noise (like extreme precip-  
 487 itation, see Figure S7), and datasets that do not cover the globe. Second, extreme pre-  
 488 cipitation has major implications for human health (Ali et al., 2019; Eekhout et al., 2018;  
 489 Rosenzweig et al., 2002) but there is considerable disagreement between climate mod-  
 490 els in its signal (Figure S7). This neural network approach can be used to identify agreed-  
 491 upon patterns despite climate model spread. Further in this section, we will demonstrate  
 492 that LRP maps can be used to investigate climate model differences and better under-  
 493 stand the time evolution of the forced response.

494 The extreme precipitation signal is not as pronounced as the temperature signal,  
 495 and using the Americas rather than the full globe limits the amount of unique informa-  
 496 tion in the input field. Nevertheless, the neural networks are still able to detect patterns  
 497 of forced change. Figure 8 depicts the years predicted by one neural network trained on  
 498 seasonal-mean extreme precipitation. As in Figure 3, the neural network is unable to ac-  
 499 curately predict the year given CMIP6 data prior to the TOE around 2010, whereafter  
 500 the predicted years generally follow the 1:1 line with the truth years, indicating that the  
 501 neural network has identified reliable indicators of change for this period. All Pearson  
 502 correlations of the actual years with the predicted years for extreme precipitation in ob-  
 503 servations are positive ( $r \approx 0.4$ ), demonstrating that the indicator patterns found in cli-  
 504 mate models can be successfully applied to observations (Figure 4). These correlations

505 are not as strong as those for mean precipitation observations, due in part to the mag-  
506 nitude of climate model disagreement in extreme precipitation as well as the observa-  
507 tional dataset used: ERA5. As shown in Figure S5, the correlations of actual with pre-  
508 dicted years for ERA5 precipitation observations are far smaller than those for GPCP  
509 observations. ERA5 tends to perform poorly in remote regions such as northern North  
510 America and northwestern South America (Bell et al., 2021), which may be responsible  
511 for these low correlations. The correlation between actual years and neural network-predicted  
512 years for extreme precipitation observations are explored in much more detail by Madakumbura  
513 et al. (2021).

514 To investigate the indicator patterns used by the neural networks to predict the  
515 year when the forced signal first emerges from the background noise, we apply LRP to  
516 all climate model samples in the training and testing sets for all 100 neural networks at  
517 the TOE (using the TOE calculated for each climate model and neural network individ-  
518 ually, see Figure S8). LRP points to western South America in DJF and British Columbia  
519 in MAM and SON as the most relevant regions when the neural networks first detect the  
520 forced response (Figure 9a-d). These LRP maps exhibit a more even distribution in rel-  
521 evance across each region and season than the end-of-the-21st-century LRP maps of global  
522 temperature and precipitation (Figures 6a-d, 7a-d). Predicting the year at the TOE, when  
523 the signal has just barely emerged from the background climate, likely requires the neu-  
524 ral networks to use all of the information available to them.

525 Up to this point, we have only considered the mean LRP maps across climate mod-  
526 els. Since the neural networks are nonlinear by nature, they can identify multiple pat-  
527 terns that differ between climate models for a given decade. We apply k-means cluster-  
528 ing to all 3200 LRP maps at the TOE (32 climate models samples, 100 neural networks)  
529 to identify two distinct indicator patterns that are being used by the climate models (Fig-  
530 ure 9e-l, see the supplementary materials for more details on k-means clustering). Tak-  
531 ing the difference between the mean LRP maps for clusters one and two reveals that the  
532 Amazon in JJA is a highly relevant region in cluster one, while western Canada in DJF  
533 is a highly relevant region in cluster two (Figure 9m-p). With the sole exception of MPI-  
534 ESM1-2-HR, all 100 LRP maps for each individual climate model fall cleanly into one  
535 cluster or the other, suggesting that there are two distinct ways in which the forced sig-  
536 nal emerges in the CMIP6 simulations (Figure 10). Interestingly, when k-means is in-  
537 structed to identify 32 unique clusters within the LRP maps, each cluster contains all  
538 100 relevance maps for each of the 32 climate models. In other words, the pathway used  
539 by the neural networks to predict the year is unique to each climate model and distin-

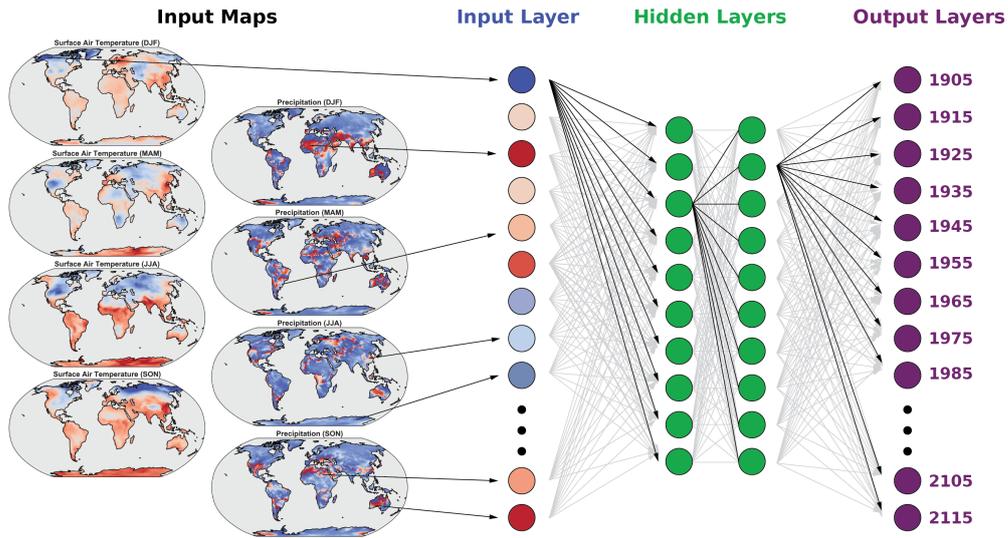
540 guishable from all other climate models, regardless of whether the climate model sam-  
541 ples appear in the training or testing sets.

542 In the same way that indicator patterns can differ between models, indicator pat-  
543 terns are also able to evolve through time (e.g., Barnes et al., 2020; Labe & Barnes, 2021;  
544 Madakumbura et al., 2021). Comparing the LRP maps at the TOE (Figure 11a-d) with  
545 those at the end of the 21st century (Figure 11e-h) highlights the regions that become  
546 more important for predicting the year over time. The difference plots in Figure 11i-l  
547 reveal that the neural network learns to focus on Alaska during MAM, JJA, and SON,  
548 Greenland in JJA and SON, and Quebec in MAM and SON as the forced response be-  
549 comes stronger. These regions are more important for predicting the year at the end of  
550 the 21st century than the early 21st century. While further exploration is required, there  
551 are several reasons a region may become more relevant over time. For example, it may  
552 be that the region does not initially have a clear forced signal, but following some abrupt  
553 change (e.g. an ice-free Arctic) the forced signal becomes extremely pronounced. It may  
554 also be that the region has a signal that is consistently agreed upon by the majority of  
555 CMIP6 climate models, and becomes more relevant compared to other regions as climate  
556 model projections in those other regions drift apart. These time-varying patterns sup-  
557 port the idea that combined indicators are effective for identifying dynamically evol-  
558 ving patterns of forced change.

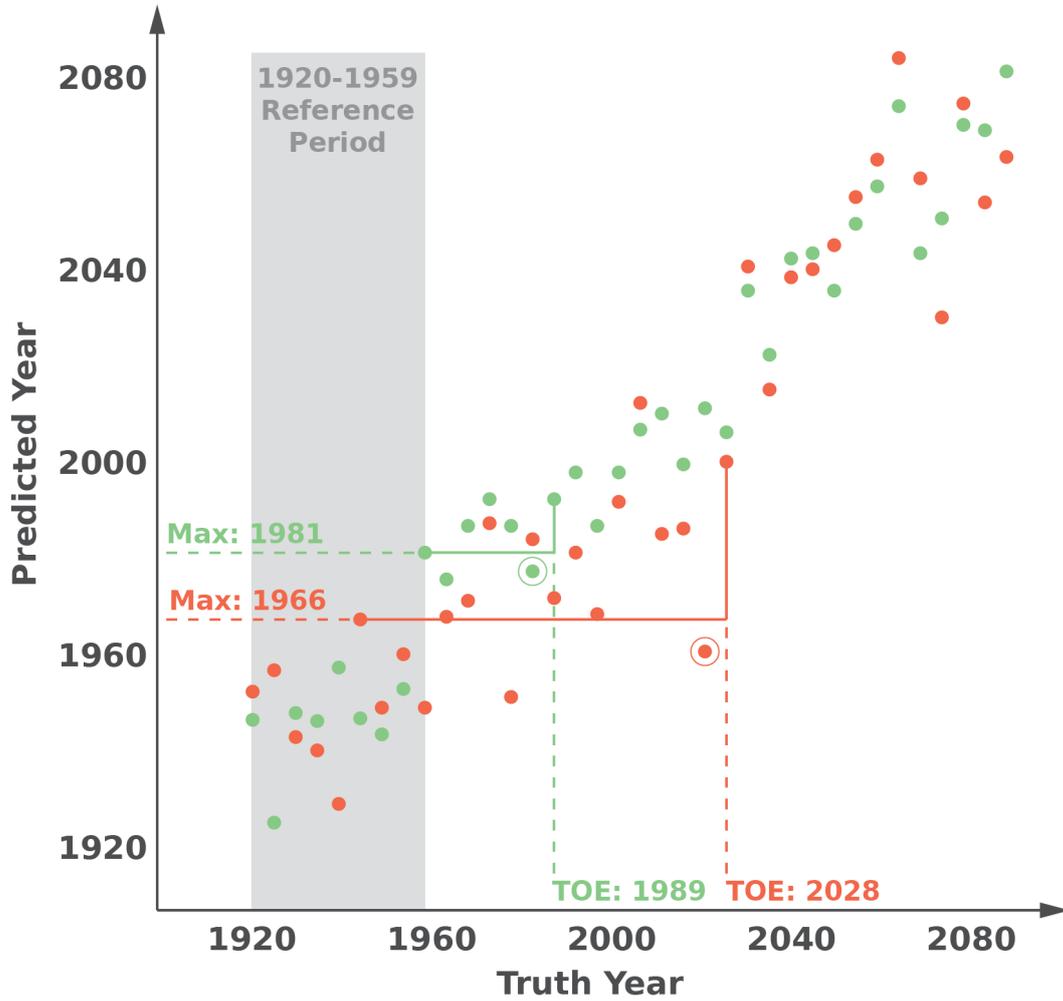
## 559 **6 Conclusions**

560 When tasked with predicting the year given climate model simulations of temper-  
561 ature, precipitation, or extreme precipitation, artificial neural networks were able to learn  
562 indicator patterns of the forced response that allow them to distinguish between maps  
563 from different years. Neural networks were trained on both single and combined fields  
564 of data, which revealed that the forced response can be detected earlier for combined fields  
565 than for single fields alone. Visualizing the decision-making process of the neural net-  
566 works with layer-wise relevance propagation revealed reliable, multivariate patterns of  
567 forced change. These indicator patterns, such as those identified for extreme precipita-  
568 tion in Section 5, vary through time and can be used to understand which regions and  
569 variables become more relevant as forced change progresses. The complex architecture  
570 of the neural networks allowed for the discovery of several different relationships between  
571 the input maps and the output decade. Two unique clusters were examined in Section  
572 5 which revealed two distinct ways in which the extreme precipitation response emerges  
573 in CMIP6 data. This work serves as an early introduction to how neural networks can  
574 be used to study combined indicators of forced change. This framework is flexible, and

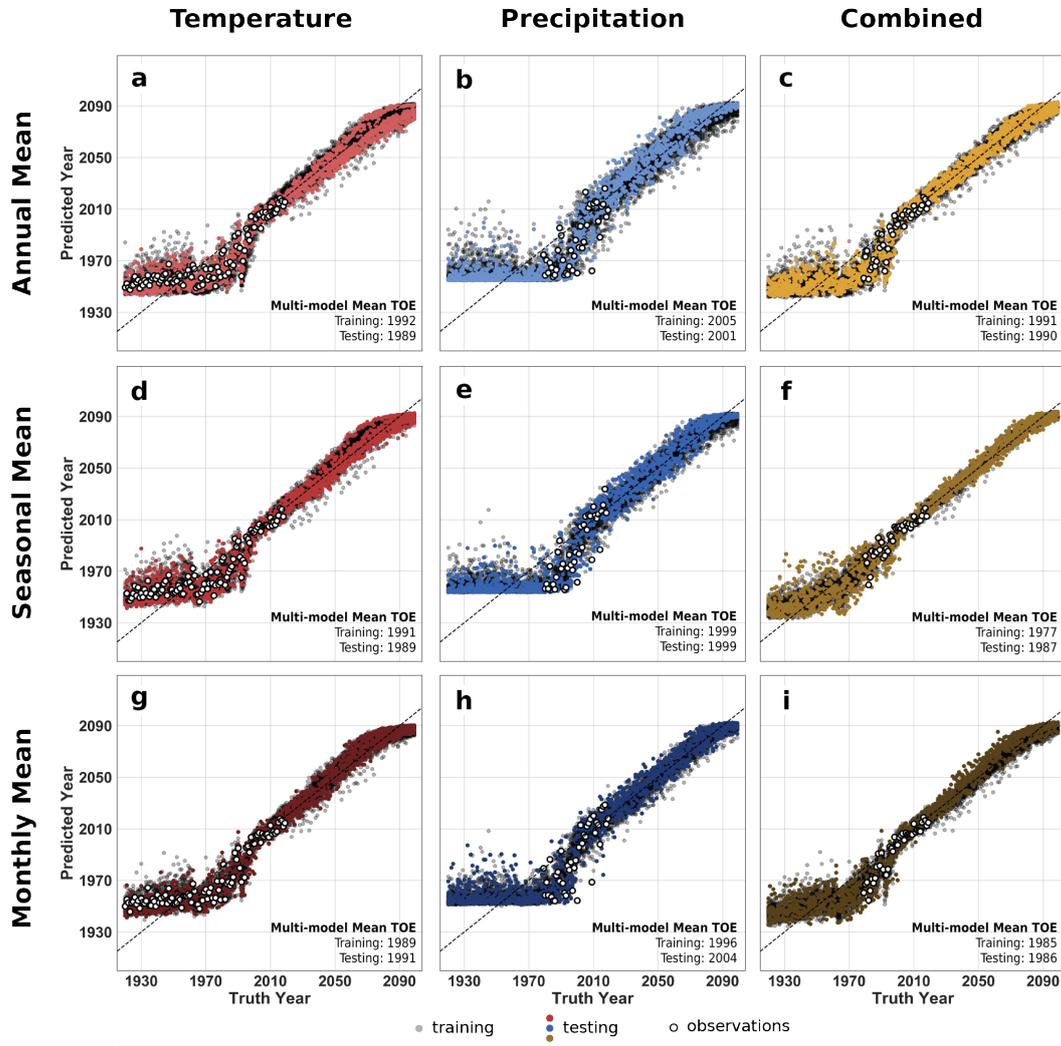
575 should be expanded to other variables, regions of focus, and climate change scenarios,  
576 to identify the combined indicators that best elucidate the forced signal. Further appli-  
577 cation of this technique to compound climate extremes, such as heat wave intensity, drought  
578 duration, and flood frequency, may reveal that explainable neural networks are useful  
579 for assessing societal impacts and improving climate change preparedness.



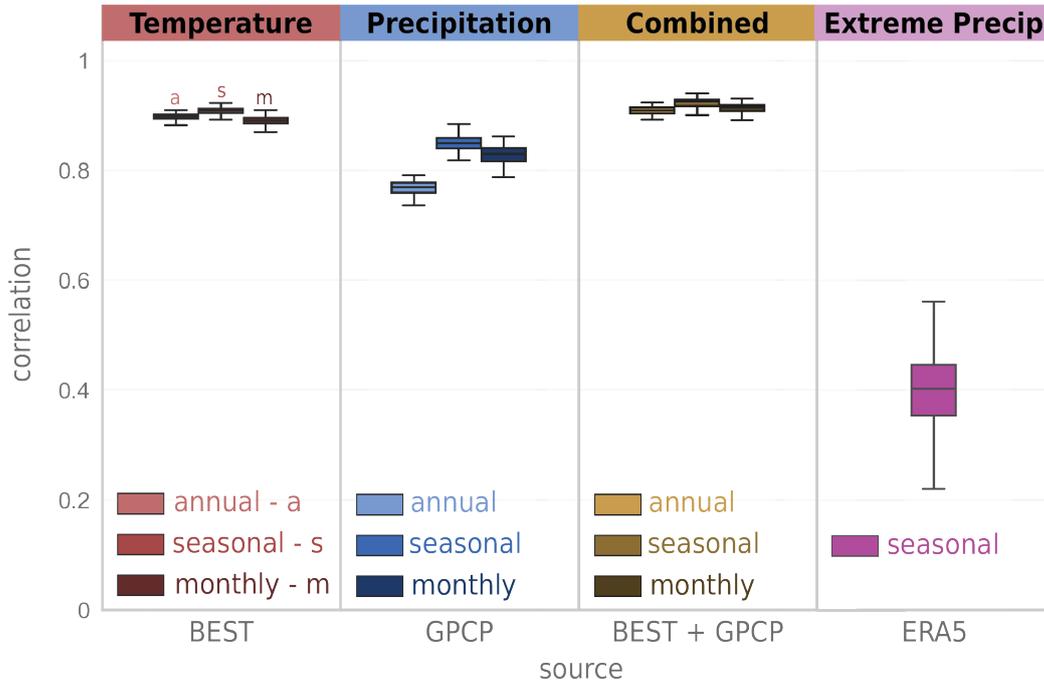
**Figure 1.** Schematic of the fully connected neural network architecture. Inputs from multiple maps of data are flattened into an input layer vector (size of the input layer ranges from 948 to 22,752). These inputs are fed through two hidden layers with ten nodes each. The neural network is trained to predict the year that the data came from, outputting the likelihood that the input data came from each decade midpoint between 1905 and 2115. This is then converted to a year via fuzzy classification.



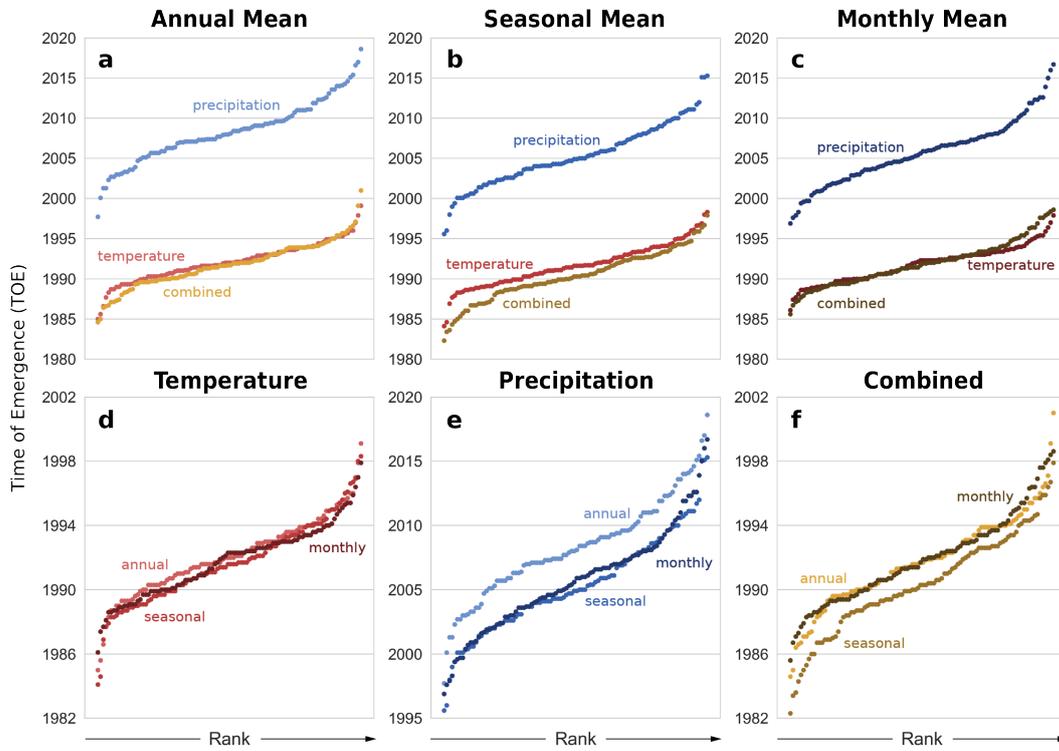
**Figure 2. Calculation of TOE.** The TOE is defined as the earliest year in which a map, and all subsequent maps, permanently exceed the maximum predicted year from the baseline period (1920-1959). The baseline maximum for each model is indicated by the horizontal lines, the last year that falls below the baseline maximum is circled, and the TOE is indicated by the vertical lines. Sample model 1 (dark red) has a baseline maximum of 1966 and permanently exceeds this threshold in 2028. Sample model 2 (light green) has a baseline maximum of 1981 and permanently exceeds this threshold in 1989. Thus, the TOE for sample models 1 and 2 are estimated as 2028 and 1989, respectively.



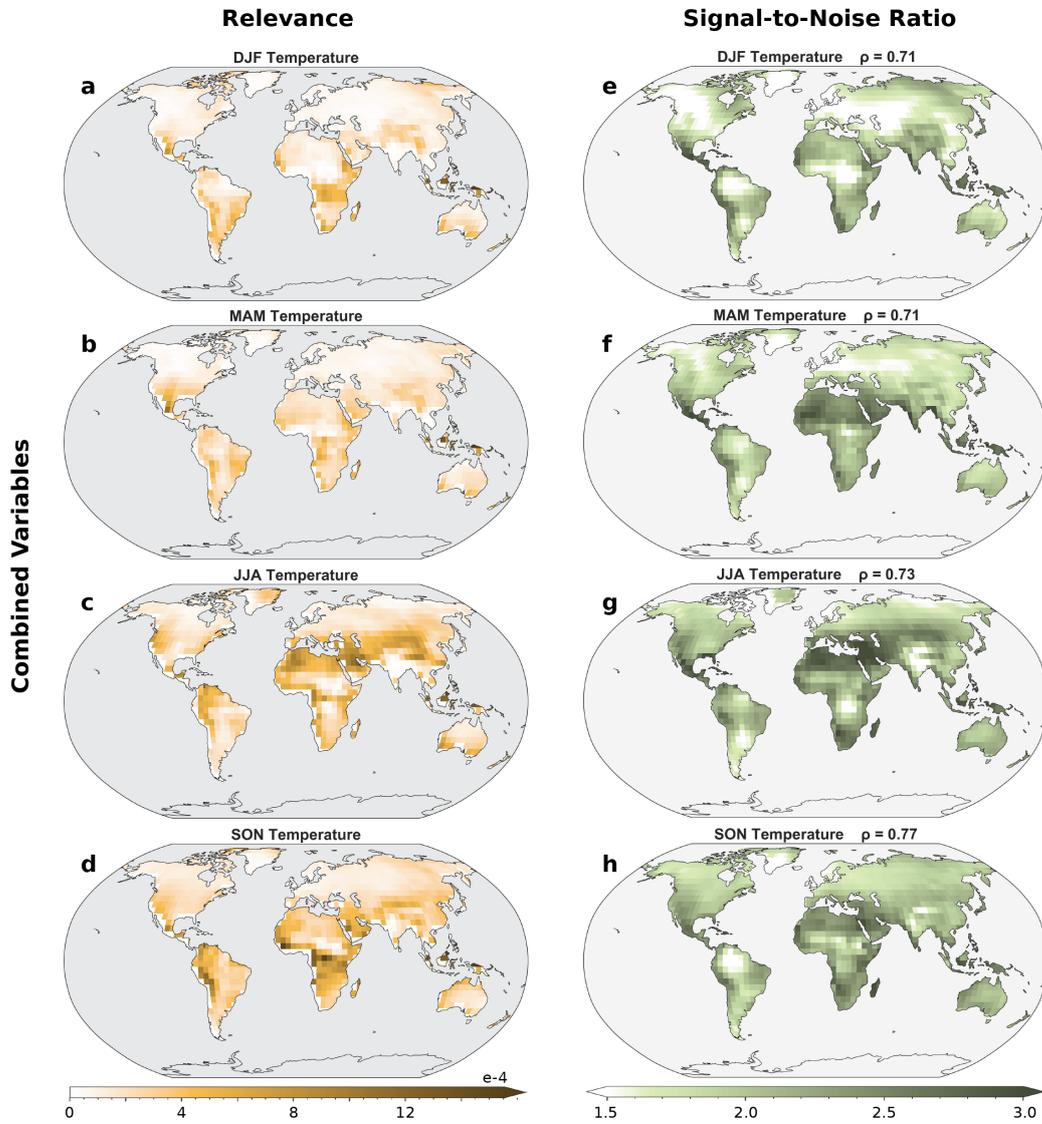
**Figure 3. Neural network output for temperature and precipitation.** Year predicted by the neural network (y-axis) versus the truth year (x-axis) for temperature (a, d, g), precipitation (b, e, h), and temperature and precipitation combined (c, f, i). Input maps include annual-mean data (a, b, c), seasonal-mean data (d, e, f), and monthly-mean data (g, h, i). Training data is shown in gray, testing data is shown in color, and observations are shown in white.



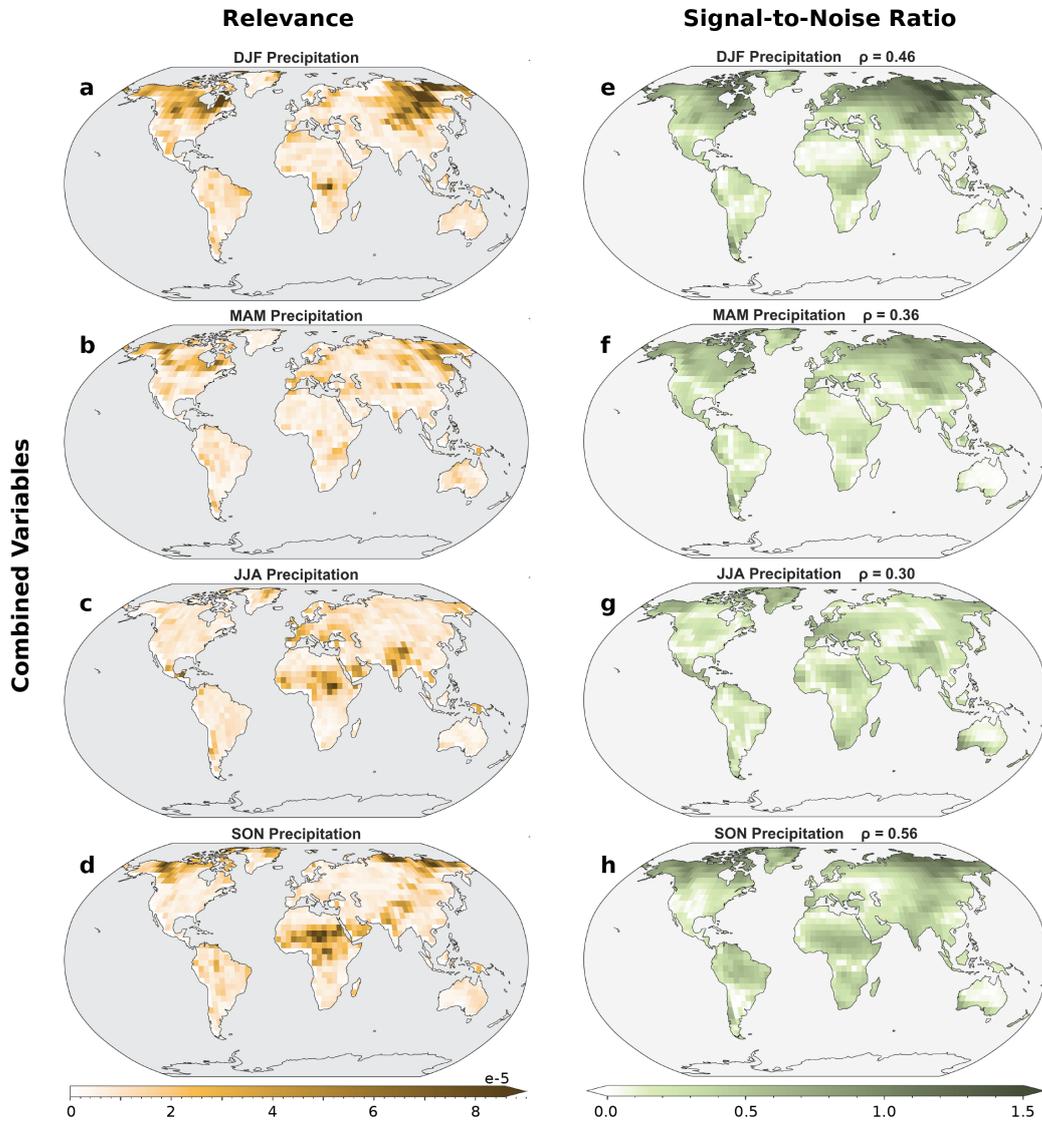
**Figure 4. Correlation of actual years with predicted years for observations.** Pearson correlations of the actual years with the years predicted by 100 trained neural networks given observations of temperature, precipitation, and extreme precipitation. Correlations were computed for all years beginning in 1980 where observational data exists for all variables. The box plots indicate the first, second, and third quartile statistics, and the whiskers denote 1.5 times the interquartile range, or the minimum/maximum value, whichever is less extreme. Outliers are excluded for clarity, but can be found in Figures S4 and S5.



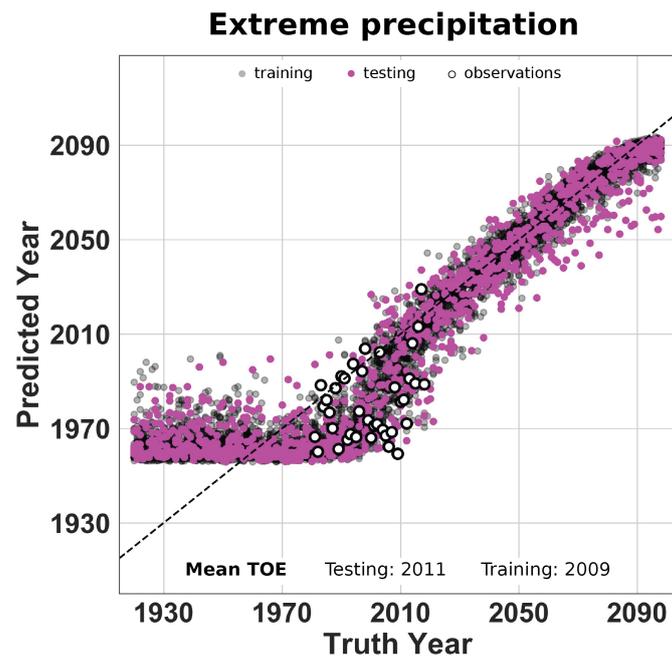
**Figure 5. Mean TOE for each input field.** Comparison of the mean time of emergence identified by neural networks trained on annual-mean (a), seasonal-mean (b), and monthly-mean (c) input fields, and neural networks trained on temperature (d), precipitation (e), and temperature and precipitation combined (f). 100 neural networks with different train-test splits were trained for each input field. Each dot represents the mean TOE for all climate models in the testing set for a single trained neural network, ranked from earliest to latest. Note the change in the y-axes between panels.



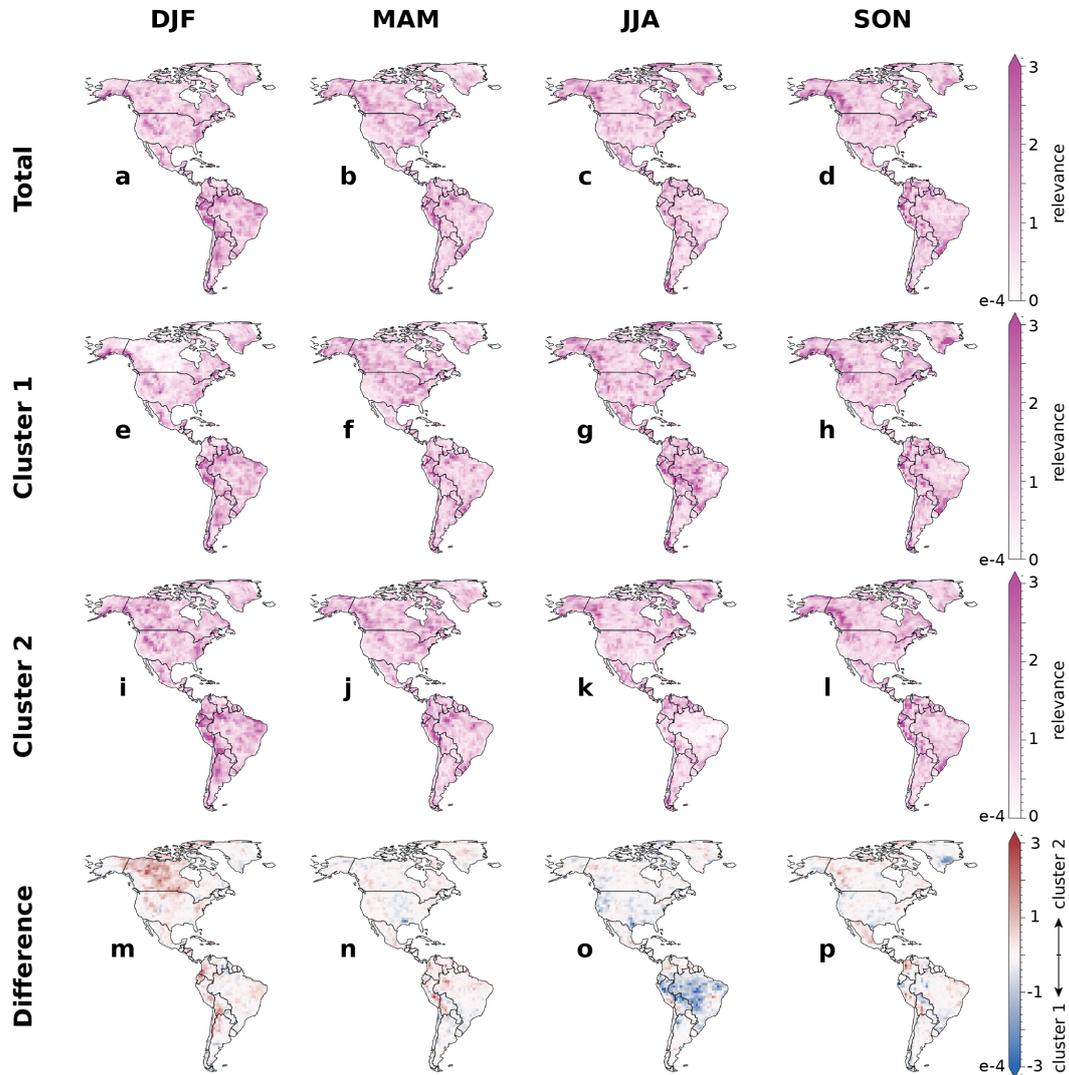
**Figure 6.** Combined indicator patterns of the forced response (temperature). Average temperature LRP results for the seasonal-mean combined neural networks (left, in yellow) and S/N ratio (right, in green) for 2090. Darker shading indicates regions of temperature that are more relevant for the neural network’s prediction or have a higher S/N ratio.



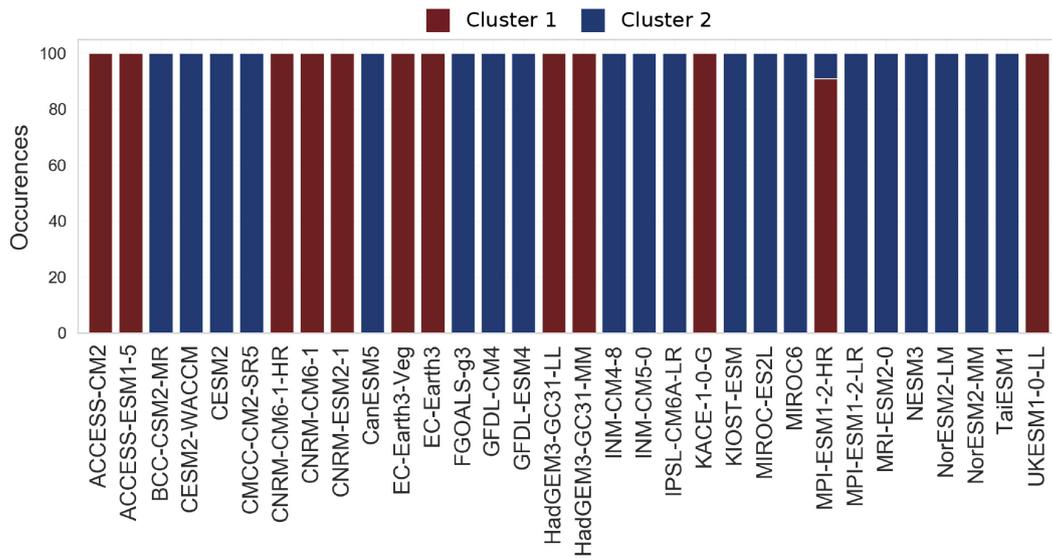
**Figure 7.** Combined indicator patterns of the forced response (precipitation). Average precipitation LRP results for the seasonal-mean combined neural networks (left, in yellow) and S/N ratio (right, in green) for 2090. Darker shading indicates regions of precipitation that are more relevant for the neural network’s prediction or have a higher S/N ratio.



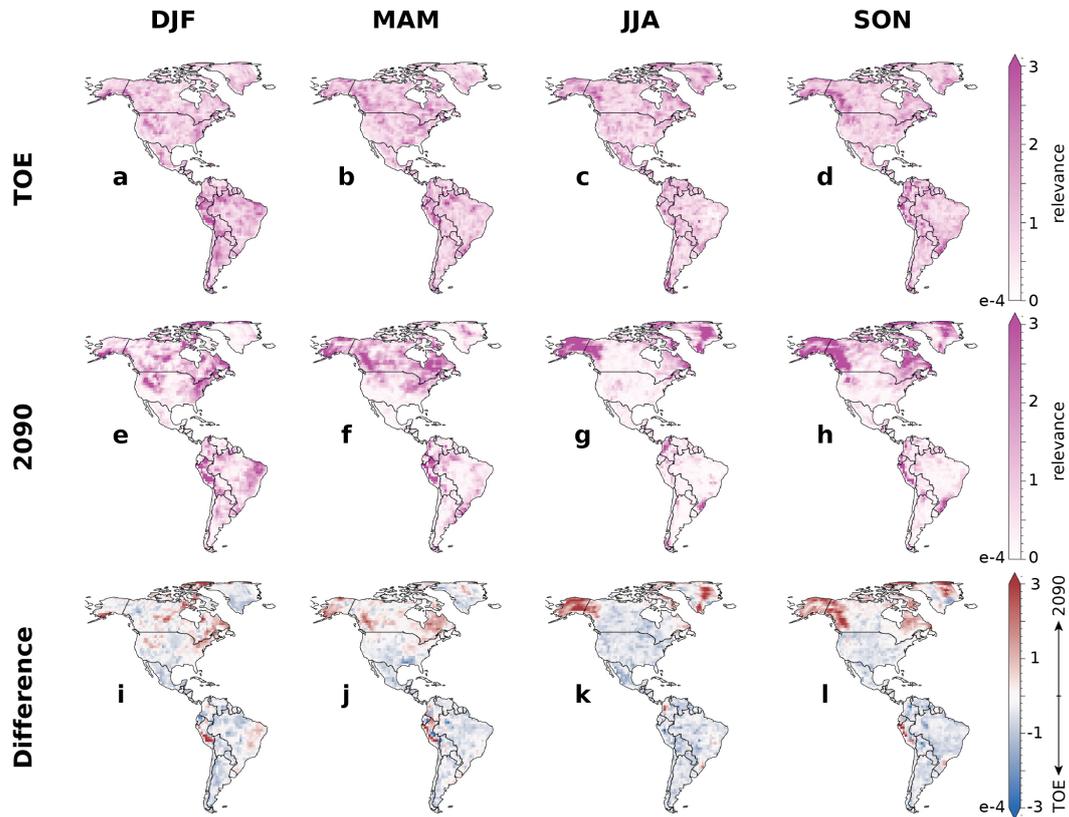
**Figure 8.** Neural network output for extreme precipitation. Year predicted by the neural network (y-axis) versus the truth year (x-axis) given seasonal-mean maps of extreme precipitation. Training data is shown in gray, testing data is shown in pink, and observations are shown in white.



**Figure 9. Relevance map clusters at the TOE for extreme precipitation.** Average LRP results for: extreme precipitation at the TOE (a-d), each cluster identified by k-means (e-h, i-l), and the difference between the clusters (m-p). In panels a-l, darker shading indicates regions of extreme precipitation that are more relevant for the neural networks' prediction of the year at the TOE. In panels m-p, blue shading indicates the regions that are more relevant in cluster 1, while red shading indicates the regions that are more relevant in cluster 2.



**Figure 10. Climate models in each relevance map cluster at the TOE.** The number of times each climate model appears in each cluster when k-means is applied to the maps of relevance at the TOE for 100 ANNs trained on extreme precipitation over the Americas. Only the relevance maps for MPI-ESM1-2-HR appear in both clusters. All other relevance maps for each climate model are found in one cluster or the other.



**Figure 11. Time evolution of extreme precipitation relevance.** Average LRP results at the TOE (a-d), 2090 (e-h), and the difference between (i-l). Darker shading in panels a-h highlights regions that were more relevant for the neural networks' prediction of the year. In panels i-l, red shading indicates regions where the relevance has increased over time, while blue shading indicates regions where the relevance has decreased over time.

## 580 **Acknowledgments**

581 I would like to thank Maria Rugenstein and Jessica Witt for their comments and feed-  
 582 back, the members of the Barnes Group for all the science discussions that brought about  
 583 new ideas, and my parents and siblings for their support through the pandemic. This  
 584 material is based upon work supported by the U.S. Department of Energy, Office of Sci-  
 585 ence, Office of Advanced Scientific Computing Research, Department of Energy Com-  
 586 putational Science Graduate Fellowship under Award Number DE-SC0020347, and by  
 587 NOAA MAPP grant NA19OAR4310289. We acknowledge the World Climate Research  
 588 Programme’s Working Group on Coupled Modelling, which is responsible for CMIP, and  
 589 we thank the climate modeling groups for producing and making available their model  
 590 output. For CMIP the U.S. Department of Energy’s Program for Climate Model Diag-  
 591 nosis and Intercomparison provides coordinating support and led development of soft-  
 592 ware infrastructure in partnership with the Global Organization for Earth System Sci-  
 593 ence Portals.

## 594 **Data Availability Statement**

595 All data used in this study is publicly available and referenced throughout the paper.  
 596 The CMIP6 simulations used in this study can be via the Earth System Grid Federa-  
 597 tion (ESGF, <https://esgf-node.llnl.gov/projects/cmip6/>). Monthly temperature obser-  
 598 vations are available through Berkeley Earth (<http://berkeleyearth.org/data/>). Global  
 599 Precipitation Climatology Project monthly global precipitation fields are available through  
 600 the NOAA Physical Sciences Laboratory (<https://psl.noaa.gov/data/gridded/data.gpcp.html>).  
 601 Monthly, daily, and sub-daily precipitation reanalyses were provided by the European Cen-  
 602 tre for Medium-Range Weather Forecasts (ERA5: [https://www.ecmwf.int/en/forecasts/datasets/reanalysis-](https://www.ecmwf.int/en/forecasts/datasets/reanalysis-datasets/era5)  
 603 [datasets/era5](https://www.ecmwf.int/en/forecasts/datasets/reanalysis-datasets/era5)) and the National Center for Atmospheric Research (JRA55: [https://climatedataguide.ucar.edu/clim-](https://climatedataguide.ucar.edu/climate-data/jra-55)  
 604 [data/jra-55](https://climatedataguide.ucar.edu/climate-data/jra-55)).

## 605 **References**

- 606 Abiodun, O. I., Jantan, A., Omolara, A. E., Dada, K. V., Mohamed, N. A., & Ar-  
 607 shad, H. (2018, November). State-of-the-art in artificial neural network appli-  
 608 cations: A survey. *Heliyon*, *4*(11), e00938. doi: 10.1016/j.heliyon.2018.e00938
- 609 Adler, R. F., Sapiano, M., Huffman, G. J., Wang, J., Gu, G., Bolvin, D., . . . Shin,  
 610 D.-B. (2018, April). The global precipitation climatology project (GPCP)  
 611 monthly analysis (new version 2.3) and a review of 2017 global precipitation.  
 612 *Atmosphere*, *9*(4). doi: 10.3390/atmos9040138
- 613 Ali, H., Modi, P., & Mishra, V. (2019, September). Increased flood risk in indian

- 614 sub-continent under the warming climate. *Weather and Climate Extremes*, 25,  
615 100212. doi: 10.1016/j.wace.2019.100212
- 616 Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.-R., & Samek, W.  
617 (2015, July). On Pixel-Wise explanations for Non-Linear classifier decisions  
618 by Layer-Wise relevance propagation. *PLoS One*, 10(7), e0130140. doi:  
619 10.1371/journal.pone.0130140
- 620 Barnes, E. A., Hurrell, J. W., Ebert-Uphoff, I., Anderson, C., & Anderson, D. (2019,  
621 November). Viewing forced climate patterns through an AI lens. *Geophys. Res.  
622 Lett.*, 46(22), 13389–13398. doi: 10.1029/2019gl084944
- 623 Barnes, E. A., Toms, B., Hurrell, J. W., Ebert-Uphoff, I., Anderson, C., & Ander-  
624 son, D. (2020, September). Indicator patterns of forced change learned by an  
625 artificial neural network. *J. Adv. Model. Earth Syst.*, 12(9), e2020MS002195.  
626 doi: 10.1029/2020ms002195
- 627 Barnett, T. P., Pierce, D. W., Hidalgo, H. G., Bonfils, C., Santer, B. D., Das, T.,  
628 ... Dettinger, M. D. (2008, February). Human-induced changes in the hy-  
629 drology of the western united states. *Science*, 319(5866), 1080–1083. doi:  
630 10.1126/science.1152538
- 631 Bell, B., Hersbach, H., Simmons, A., Berrisford, P., Dahlgren, P., Horányi, A., ...  
632 Thépaut, J. (2021). The ERA5 global reanalysis: Preliminary extension to  
633 1950. *Quart J Roy Meteor Soc.*
- 634 Bindoff, N. L., Stott, P. A., AchutaRao, K. M., Allen, M. R., Gillett, N., Gutzler,  
635 D., ... Zhang, X. (2013). Chapter 10 - detection and attribution of climate  
636 change: From global to regional. In *Climate change 2013: The physical sci-  
637 ence basis. IPCC working group I contribution to AR5.* Cambridge, United  
638 Kingdom: Cambridge University Press.
- 639 Bonfils, C. J. W., Santer, B. D., Fyfe, J. C., Marvel, K., Phillips, T. J., & Zimmer-  
640 man, S. R. H. (2020, July). Human influence on joint changes in temperature,  
641 rainfall and continental aridity. *Nat. Clim. Chang.*, 10(8), 726–731. doi:  
642 10.1038/s41558-020-0821-1
- 643 Brenowitz, N. D., & Bretherton, C. S. (2018, June). Prognostic validation of a  
644 neural network unified physics parameterization. *Geophys. Res. Lett.*, 45(12),  
645 6289–6298. doi: 10.1029/2018gl078510
- 646 Cui, L., Wang, L., Qu, S., Singh, R. P., Lai, Z., & Yao, R. (2019, April). Spa-  
647 tiotemporal extremes of temperature and precipitation during 1960–2015 in  
648 the yangtze river basin (china) and impacts on vegetation dynamics. *Theor.  
649 Appl. Climatol.*, 136(1), 675–692. doi: 10.1007/s00704-018-2519-0
- 650 Dai, A., Lin, X., & Hsu, K.-L. (2007, October). The frequency, intensity, and di-

- 651       urnal cycle of precipitation in surface and satellite observations over low- and  
652       mid-latitudes. *Clim. Dyn.*, 29(7-8), 727–744. doi: 10.1007/s00382-007-0260-y
- 653       Donat, M. G., Alexander, L. V., Herold, N., & Dittus, A. J. (2016, October). Tem-  
654       perature and precipitation extremes in century-long gridded observations,  
655       reanalyses, and atmospheric model simulations. *J. Geophys. Res.*, 121(19),  
656       11,174–11,189. doi: 10.1002/2016jd025480
- 657       Eekhout, J. P. C., Hunink, J. E., Terink, W., & de Vente, J. (2018, April).  
658       Why increased extreme precipitation under climate change negatively af-  
659       fects water security. *Hydrol. Earth Syst. Sci. Discuss.*, 22(11), 1–16. doi:  
660       10.5194/hess-2018-161
- 661       Eyring, V., Bony, S., Meehl, G. A., Senior, C. A., Stevens, B., Stouffer, R. J., &  
662       Taylor, K. E. (2016, May). Overview of the coupled model intercomparison  
663       project phase 6 (CMIP6) experimental design and organization. *Geosci. Model*  
664       *Dev.*, 9(5), 1937–1958. doi: 10.5194/gmd-9-1937-2016
- 665       Field, C. B., Barros, V. R., Mastrandrea, M. D., Mach, K. J., Abdrabo, M. A.-K.,  
666       Adger, N., . . . Yohe, G. W. (2014). Summary for policymakers. In *Climate*  
667       *change 2014: Impacts, adaptation, and vulnerability. part a: Global and sec-*  
668       *toral aspects. contribution of working group II to the fifth assessment report of*  
669       *the intergovernmental panel on climate change* (pp. 1–32). Cambridge, United  
670       Kingdom and New York, NY, USA: Cambridge University Press.
- 671       Fischer, E. M., & Knutti, R. (2012, September). Robust projections of combined hu-  
672       midity and temperature extremes. *Nat. Clim. Chang.*, 3(2), 126–130. doi: 10  
673       .1038/nclimate1682
- 674       Gaetani, M., Janicot, S., Vrac, M., Famien, A. M., & Sultan, B. (2020, May). Ro-  
675       bust assessment of the time of emergence of precipitation change in west africa.  
676       *Sci. Rep.*, 10(1), 7670. doi: 10.1038/s41598-020-63782-2
- 677       Gettelman, A., Gagne, D. J., Chen, C.-C., Christensen, M. W., Lebo, Z. J., Mor-  
678       rison, H., & Gantos, G. (2021, February). Machine learning the warm  
679       rain process. *J. Adv. Model. Earth Syst.*, 13(2), e2020MS002268. doi:  
680       10.1029/2020ms002268
- 681       Hawkins, E., & Sutton, R. (2009, August). The potential to narrow uncertainty in  
682       regional climate predictions. *Bull. Am. Meteorol. Soc.*, 90(8), 1095–1108. doi:  
683       10.1175/2009BAMS2607.1
- 684       Hawkins, E., & Sutton, R. (2012, January). Time of emergence of climate signals.  
685       *Geophys. Res. Lett.*, 39(1). doi: 10.1029/2011gl050087
- 686       Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J.,  
687       . . . Jean-Noël Thépaut (2020, July). The ERA5 global reanalysis. *Quart. J.*

- 688 *Roy. Meteor. Soc.*, 146(730), 1999–2049. doi: 10.1002/qj.3803
- 689 Kim, Y.-H., Min, S.-K., Zhang, X., Sillmann, J., & Sandstad, M. (2020, September).  
 690 Evaluation of the CMIP6 multi-model ensemble for climate extreme indices.  
 691 *Weather and Climate Extremes*, 29, 100269. doi: 10.1016/j.wace.2020.100269
- 692 Labe, Z. M., & Barnes, E. A. (2021, June). Detecting climate signals using explain-  
 693 able AI with single-forcing large ensembles. *J. Adv. Model. Earth Syst.*, 13(6),  
 694 e2021MS002464. doi: 10.1029/2021ms002464
- 695 Lagerquist, R., McGovern, A., & Gagne, D. J., II. (2019, August). Deep learning for  
 696 spatially explicit prediction of Synoptic-Scale fronts. *Weather Forecast.*, 34(4),  
 697 1137–1160. doi: 10.1175/WAF-D-18-0183.1
- 698 Lapuschkin, S. (2019). *Opening the machine learning black box with layer-wise*  
 699 *relevance propagation* (Doctoral dissertation, Technischen Universität Berlin,  
 700 Berlin, Germany). doi: 10.14279/DEPOSITONCE-7942
- 701 Lee, Y., Kummerow, C. D., & Ebert-Uphoff, I. (2021, April). Applying machine  
 702 learning methods to detect convection using geostationary operational environ-  
 703 mental satellite-16 (GOES-16) advanced baseline imager (ABI) data. *Atmos.*  
 704 *Meas. Tech.*, 14(4), 2699–2716. doi: 10.5194/amt-14-2699-2021
- 705 Li, J., Thompson, D. W. J., Barnes, E. A., & Solomon, S. (2017, December). Quan-  
 706 tifying the lead time required for a linear trend to emerge from natural climate  
 707 variability. *J. Clim.*, 30(24), 10179–10191. doi: 10.1175/JCLI-D-16-0280.1
- 708 Madakumbura, G. D., Thackeray, C. W., Norris, J., Goldenson, N., & Hall, A.  
 709 (2021, July). Anthropogenic influence on extreme precipitation over global  
 710 land areas seen in multiple observational datasets. *Nat. Commun.*, 12(1), 3944.  
 711 doi: 10.1038/s41467-021-24262-x
- 712 Maher, N., Milinski, S., & Ludwig, R. (2021). Large ensemble climate model simula-  
 713 tions: introduction, overview, and future prospects for utilising multiple types  
 714 of large ensemble. *Earth System Dynamics*, 12(2), 401–418.
- 715 Mahony, C. R., & Cannon, A. J. (2018, February). Wetter summers can intensify  
 716 departures from natural variability in a warming climate. *Nat. Commun.*, 9(1),  
 717 783. doi: 10.1038/s41467-018-03132-z
- 718 Mamalakis, A., Ebert-Uphoff, I., & Barnes, E. A. (2021). Neural network attribu-  
 719 tion methods for problems in geoscience: A novel synthetic benchmark dataset.  
 720 *arXiv*. Retrieved from <http://arxiv.org/abs/2103.10005>
- 721 Mankin, J. S., Lehner, F., Coats, S., & McKinnon, K. A. (2020, October). The  
 722 value of initial condition large ensembles to robust adaptation decision-making.  
 723 *Earths Future*, 8(10). doi: 10.1029/2020ef001610
- 724 Marvel, K., & Bonfils, C. (2013, November). Identifying external influences on global

- 725 precipitation. *Proc. Natl. Acad. Sci. U. S. A.*, *110*(48), 19301–19306. doi: 10  
726 .1073/pnas.1314382110
- 727 Montavon, G., Binder, A., Lapuschkin, S., Samek, W., & Müller, K.-R. (2019).  
728 Layer-Wise relevance propagation: An overview. In W. Samek, G. Montavon,  
729 A. Vedaldi, L. K. Hansen, & K.-R. Müller (Eds.), *Explainable AI: Interpreting,*  
730 *explaining and visualizing deep learning* (Vol. 11700, pp. 193–209). Cham:  
731 Springer International Publishing. doi: 10.1007/978-3-030-28954-6\\_10
- 732 Mudelsee, M. (2019, March). Trend analysis of climate time series: A review of  
733 methods. *Earth-Sci. Rev.*, *190*, 310–322. doi: 10.1016/j.earscirev.2018.12.005
- 734 North, G. R., & Stevens, M. J. (1998, April). Detecting climate signals in  
735 the surface temperature record. *J. Clim.*, *11*(4), 563–577. doi: 10.1175/  
736 1520-0442(1998)011<0563:DCSITS>2.0.CO;2
- 737 O’Neill, B. C., Tebaldi, C., van Vuuren, D. P., Eyring, V., Friedlingstein, P., Hurtt,  
738 G., ... Sanderson, B. M. (2016, September). The scenario model intercompari-  
739 son project (ScenarioMIP) for CMIP6. *Geoscientific Model Development*, *9*(9),  
740 3461–3482. doi: 10.5194/gmd-9-3461-2016
- 741 Rohde, R. A., & Hausfather, Z. (2020, December). The berkeley earth land/ocean  
742 temperature record. *Earth Syst. Sci. Data*, *12*(4), 3469–3479. doi: 10.5194/  
743 essd-12-3469-2020
- 744 Rosenzweig, C., Tubiello, F. N., Goldberg, R., Mills, E., & Bloomfield, J. (2002, Oc-  
745 tober). Increased crop damage in the US from excess precipitation under cli-  
746 mate change. *Glob. Environ. Change*, *12*(3), 197–202. doi: 10.1016/S0959  
747 -3780(02)00008-0
- 748 Sanderson, B. M., Oleson, K. W., Strand, W. G., Lehner, F., & O’Neill, B. C. (2018,  
749 February). A new ensemble of GCM simulations to assess avoided impacts in  
750 a climate mitigation scenario. *Clim. Change*, *146*(3), 303–318. doi: 10.1007/  
751 s10584-015-1567-z
- 752 Santer, B. D., Fyfe, J. C., Solomon, S., Painter, J. F., Bonfils, C., Pallotta, G., &  
753 Zelinka, M. D. (2019, October). Quantifying stochastic uncertainty in detec-  
754 tion time of human-caused climate signals. *Proc. Natl. Acad. Sci. U. S. A.*,  
755 *116*(40), 19821–19827. doi: 10.1073/pnas.1904586116
- 756 Santer, B. D., Mears, C., Doutriaux, C., Caldwell, P., Gleckler, P. J., Wigley,  
757 T. M. L., ... Wentz, F. J. (2011, November). Separating signal and noise  
758 in atmospheric temperature changes: The importance of timescale. *J. Geophys.*  
759 *Res.*, *116*(D22). doi: 10.1029/2011jd016263
- 760 Santer, B. D., Taylor, K. E., Wigley, T. M. L., Johns, T. C., Jones, P. D., Karoly,  
761 D. J., ... Tett, S. (1996, July). A search for human influences on the

- 762 thermal structure of the atmosphere. *Nature*, *382*(6586), 39–46. doi:  
 763 10.1038/382039a0
- 764 Scaife, A. A., & Smith, D. (2018, July). A signal-to-noise paradox in climate science.  
 765 *npj Climate and Atmospheric Science*, *1*(1), 1–8. doi: 10.1038/s41612-018-0038  
 766 -4
- 767 Schneider, T., & Held, I. M. (2001, February). Discriminants of Twentieth-Century  
 768 changes in earth surface temperatures. *J. Clim.*, *14*(3), 249–254. doi: 10.1175/  
 769 1520-0442(2001)014<0249:LDOTCC>2.0.CO;2
- 770 Schulzweida, U. (2019). CDO user guide (version 1.9. 6). *Max Planck Institute for*  
 771 *Meteorology: Hamburg, Germany.*
- 772 Silva, S. J., Ma, P.-L., Hardin, J. C., & Rothenberg, D. (2021, May). Physically reg-  
 773 ularized machine learning emulators of aerosol activation. *Geosci. Model Dev.*,  
 774 *14*(5), 3067–3077. doi: 10.5194/gmd-14-3067-2021
- 775 Sippel, S., Meinshausen, N., Fischer, E. M., Székely, E., & Knutti, R. (2020, Jan-  
 776 uary). Climate change now detectable from any single day of weather at global  
 777 scale. *Nat. Clim. Chang.*, *10*(1), 35–41. doi: 10.1038/s41558-019-0666-7
- 778 Solomon, A., & Newman, M. (2012, July). Reconciling disparate twentieth-century  
 779 Indo-Pacific ocean temperature trends in the instrumental record. *Nat. Clim.*  
 780 *Chang.*, *2*(9), 691–699. doi: 10.1038/nclimate1591
- 781 Solow, A. R. (1987, October). Testing for climate change: An application of the  
 782 Two-Phase regression model. *J. Appl. Meteorol. Climatol.*, *26*(10), 1401–1405.  
 783 doi: 10.1175/1520-0450(1987)026<1401:TFCCAA>2.0.CO;2
- 784 Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R.  
 785 (2014). Dropout: a simple way to prevent neural networks from overfitting.  
 786 *J. Mach. Learn. Res.*, *15*(1), 1929–1958.
- 787 Tabari, H., & Willems, P. (2018, March). Seasonally varying footprint of climate  
 788 change on precipitation in the middle east. *Sci. Rep.*, *8*(1), 4435. doi: 10.1038/  
 789 s41598-018-22795-8
- 790 Toms, B. A., Barnes, E. A., & Ebert-Uphoff, I. (2020, September). Physically  
 791 interpretable neural networks for the geosciences: Applications to earth sys-  
 792 tem variability. *J. Adv. Model. Earth Syst.*, *12*(9), e2019MS002002. doi:  
 793 10.1029/2019ms002002
- 794 USGCRP. (2018). *Impacts, risks, and adaptation in the united states: Fourth na-*  
 795 *tional climate assessment, volume II* (Tech. Rep.). Washington, DC, USA:  
 796 U.S. Global Change Research Program. doi: 10.7930/NCA4.2018
- 797 Weyn, J. A., Durran, D. R., & Caruana, R. (2020, September). Improving data-  
 798 driven global weather prediction using deep convolutional neural networks

- 799 on a cubed sphere. *J. Adv. Model. Earth Syst.*, *12*(9), e2020MS002109. doi:  
800 10.1029/2020ms002109
- 801 Wills, R. C., Battisti, D. S., Armour, K. C., Schneider, T., & Deser, C. (2020, Oc-  
802 tober). Pattern recognition methods to separate forced responses from inter-  
803 nal variability in climate model ensembles and observations. *J. Clim.*, *33*(20),  
804 8693–8719. doi: 10.1175/JCLI-D-19-0855.1
- 805 Wills, R. C., Schneider, T., Wallace, J. M., Battisti, D. S., & Hartmann, D. L.  
806 (2018, March). Disentangling global warming, multidecadal variability, and  
807 el niño in pacific temperatures. *Geophys. Res. Lett.*, *45*(5), 2487–2496. doi:  
808 10.1002/2017GL076327
- 809 Zadeh, L. A. (1965). Information and control. *Fuzzy Sets and Systems*, *8*(3), 338–  
810 353. doi: 10.1002/joc.1027
- 811 Zappa, G., Hoskins, B. J., & Shepherd, T. G. (2015, August). Improving climate  
812 change detection through optimal seasonal averaging: The case of the north  
813 atlantic jet and european precipitation. *J. Clim.*, *28*(16), 6381–6397. doi:  
814 10.1175/JCLI-D-14-00823.1