

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

Innovative methods are required to determine which behaviors of the climate are the result of the forced signal and which are the result of climate noise. Decades of research have provided a diverse toolkit for this task (North & Stevens, 1998) which includes linear regression (e.g., Mudelsee, 2019; Santer et al., 1996; Sippel et al., 2020; Solow, 1987), empirical orthogonal functions and linear discriminant analysis (e.g., Santer et al., 2019; Schneider & Held, 2001; Wills et al., 2018, 2020), and linear inverse models (e.g., Solomon & Newman, 2012), to name a few. Recently, neural networks have also entered the fold. Neural networks are machine learning algorithms that are able to detect complex, nonlinear relationships between input and output data (Abiodun et al., 2018). Because neural networks are able to detect highly complex relationships, they are useful for many high dimensional problems and have become prevalent in several atmospheric science research fields, such as weather forecasting (e.g., Lagerquist et al., 2019; Lee et al., 2021; Weyn et al., 2020), climate model parameterizations (e.g., Brenowitz & Bretherton, 2018; Gettelman et al., 2021; Silva et al., 2021), and, most relevant to the focus of this study, detection of a forced climate response (e.g., Barnes et al., 2019, 2020; Labe & Barnes, 2021; Madakumbura et al., 2021). To detect patterns of forced change, Barnes et al. (2020) trained a neural network to predict the year label of maps of annual-mean temperature (or precipitation) from climate model simulations for forced historical and future scenarios. Given that the internal variability in any given year differs between the various climate models, the neural network had to learn patterns of the forced climate response. Using neural network explainability methods, they then visualized the regions that were most reliable indicators for identifying change across the CMIP5 models. Barnes et al. (2020) demonstrated that neural networks, and their explainability methods, are powerful tools for extracting forced patterns from climate data. This neural network method is a natural approach for isolating the forced climate response. While many other methods require assumptions to be made about the time evolution of the forced signal and internal variability within the system, neural networks do not (Barnes et al., 2019). Following Barnes et al. (2020), neural networks have since been used to explore the sensitivity of regional temperature signals to aerosols and greenhouse gases using single-forcing

87 large ensembles, and to detect the signal of extreme precipitation in observational datasets  
88 (Labe & Barnes, 2021; Madakumbura et al., 2021).

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

113         Combining fields can be useful for identifying patterns of forced change that do not  
114 reveal themselves in single fields alone, but this added information does not come with-  
115 out its drawbacks. Many variables covary in complex and nonlinear ways, such as sea  
116 surface temperature and precipitation (Lu et al., 2015), drought indices (Wu et al., 2017),  
117 and snowpack, soil moisture and flood risk (Swain et al., 2020), often requiring complex  
118 statistics to isolate these interactions. Identifying nonlinear correlations within climate  
119 fields introduces another issue, namely in explaining the complex interplay between fields.  
120 These drawbacks highlight the need for methods that are both complex and explainable  
121 in multivariate climate analyses.

122 Providing a method for both nonlinear and multi-variable analysis of the forced re-  
 123 sponse, this study extends the neural-network approach of Barnes et al. (2020) to com-  
 124 bined fields of input. Combined fields could mean the same variable for different tem-  
 125 poral segments (e.g. seasons), or different geophysical variables, both of which are ex-  
 126 plored here. For the sake of consistency and comparability, this study largely follows the  
 127 methodology of Barnes et al. (2020), however there are some departures. We standard-  
 128 ize the input fields differently which improves the predictive skill of the neural networks.  
 129 We also use a slightly simpler neural network architecture to reduce the computational  
 130 expense of training a single neural network, and the results from multiple neural networks,  
 131 rather than just one, are explored. Barnes et al. (2020) demonstrated the utility of neu-  
 132 ral network explainability methods, and we use these methods in tandem with a clus-  
 133 tering technique to enhance post-hoc explanations of neural network decisions.

134 Section 2 outlines the climate models and observations analyzed in this study. Sec-  
 135 tion 3 introduces the neural network design, the explainability technique (layer-wise rel-  
 136 evance propagation; LRP), and their applications to detection of the forced climate re-  
 137 sponse. We then apply these methods to global temperature and precipitation over land  
 138 in Section 4. Here we investigate the benefits of combining variables and compare the  
 139 results of the neural network with the classical approach of calculating signal-to-noise  
 140 ratios. In Section 5, we explore the patterns of the forced response for extreme precip-  
 141 itation over the Americas and investigate the applications of LRP to studying the evo-  
 142 lution of nonlinear climate patterns across multiple climate models. Finally, Section 6  
 143 summarizes the results of this work and its implications for future work in forced change  
 144 detection.

## 145 **2 Data**

### 146 **2.1 CMIP6 Climate Models**

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

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

165 The neural networks are trained on CMIP6 climate model data. One ensemble mem-  
166 ber is selected for each of the 37 CMIP6 climate models analyzed so each climate model  
167 is only represented once in the training and testing data. Since daily output is required  
168 to calculate very wet days, we are limited to 32 models for extreme precipitation (Fig-  
169 ure S3). We analyze the climate model data from 1920 to 2098 under historical forcing  
170 (1920–2014) and the SSP585 scenario (2015–2098). SSP585 represents the highest de-  
171 velopment pathway within CMIP6 scenarios (O’Neill et al., 2016), combining shared so-  
172 cioeconomic pathway 5 (SSP5) and representative concentration pathway 8.5 (RCP8.5).

173 Our neural network methodology requires that all climate model fields have the same  
174 shape. To accommodate this we regrid the climate model fields from their native res-  
175 olutions using the second-order conservative remapping method in the Climate Data Op-  
176 erators package from MPI (Schulzweida, 2019). This regridding step reduces the spatial  
177 resolution of the data for most climate models. For temperature and precipitation, the  
178 data is regridded to 4 degrees latitude by 4 degrees longitude. We elect to use lower res-  
179 olution data to reduce the computational expense of training neural networks over global  
180 maps of temperature and precipitation. Since the domain for extreme precipitation is  
181 smaller than the domain for temperature and precipitation (see the following paragraph),  
182 and higher resolution data may better capture regional extreme precipitation patterns,  
183 the data for extreme precipitation is regridded to a slightly higher resolution: 1.5 degrees  
184 latitude by 1.5 degrees longitude.

185 Two spatial domains are considered in the results of this paper. For temperature  
186 and precipitation, the neural networks are trained on all land north of 60°S. Here, we  
187 choose to focus on land grid points because that is where humanity lives and will acutely  
188 feel the impacts of changing surface temperatures and precipitation. We also exclude Antarc-  
189 tica where climate models and reanalyses struggle to accurately simulate temperature  
190 and precipitation. Each map of temperature and precipitation has 948 unique data points.  
191 For extreme precipitation, the neural networks are trained on North and South Amer-

192 ica (land grid points bounded by 90°N, 55°S, 170°W, and 25°W). Here, we choose to nar-  
 193 row the regional scope to show that neural networks are powerful tools for identifying  
 194 the forced response even when the spatial domain, and thus the available data, is lim-  
 195 ited. Each map of extreme precipitation has 2314 unique data points.

## 196 2.2 Observations

197 While this work largely focuses on the results of neural networks trained and tested  
 198 on climate model data, we show that neural networks trained on climate model data can  
 199 be applied to observational data as well. For temperature, we use the Berkeley Earth  
 200 Surface Temperature (BEST) dataset (Rohde & Hausfather, 2020). This dataset pro-  
 201 vides both a temperature climatology and the anomalies at monthly resolution from 1850  
 202 to the present. We added the anomalies to the climatology to reconstruct the absolute  
 203 temperature ( $K$ ) at each grid point for all months between 1920 and 2019. Monthly ob-  
 204 servational precipitation fields are obtained from the NOAA Global Precipitation Cli-  
 205 matology Project (GPCP), version 2.3, for 1979 to the present (Adler et al., 2018). Since  
 206 daily precipitation fields are required to calculate extreme precipitation, and daily GPCP  
 207 precipitation observations are only available back to October 1996, we elected to calcu-  
 208 late observed extreme precipitation using the European Centre for Medium-Range Weather  
 209 Forecasts' ERA5 global reanalysis (Hersbach et al., 2020) at 6-hour resolution from 1980  
 210 to present. All observations are regridded in the same way as the climate model data for  
 211 each respective variable.

## 212 3 Forced Change Detection Framework

### 213 3.1 Neural Network Design

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

223 The neural network architecture is illustrated in Figure 1. Each unit of the input  
 224 layer corresponds to a different grid point in the input fields. For example, a neural net-

225 work that uses seasonal-mean maps of temperature and precipitation as input (two vari-  
226 ables and four seasons for a total of eight maps, 948 grid points per map) would have  
227 an input vector with 7,584 units. In all cases, this input layer is followed by two fully  
228 connected hidden layers with ten nodes each. The hidden layers are followed by an out-  
229 put layer that consists of 22 classes, one corresponding to each decade midpoint between  
230 1905 and 2115 (e.g. 1905, 1915, 1925, . . . , 2115). A softmax function is applied to the  
231 outputs to convert them to units of likelihood, where the sum of the output vector is one.

232 This is a relatively simple architecture for a neural network. Neural networks with  
233 this architecture learn the patterns of forced change well, and more complicated archi-  
234 tectures do not substantially improve neural network performance (see Figure S2). It is  
235 also notable that this neural network architecture performs better than multiple linear  
236 regression, especially when trained on precipitation, and thus using nonlinear techniques  
237 improves our ability to detect the year via patterns of forced change (Figure S2). This  
238 simple architecture is also widely accessible to most in the climate science community  
239 as it can be trained on a personal laptop—highly complex architectures can be prohibitively  
240 computationally expensive (Chen et al., 2020). These neural networks were trained on  
241 a standard desktop computer with 16 GB of RAM and a 3.1 GHz, 6-core processor. Train-  
242 ing a single network took anywhere between two and ten minutes depending on the size  
243 of the input field. More details on the neural network design and hyperparameter tun-  
244 ing can be found in the supplementary materials.

245 The neural network is tasked with “predicting the year” rather than “predicting  
246 the decade” as the output layer may suggest. To translate between decade midpoints and  
247 individual year labels, we use fuzzy encoding (Zadeh, 1965) such that each year can be  
248 mapped to one or more neighboring classes with varying degrees of membership (encoded  
249 as likelihood). This is different than traditional methods that would map each year to  
250 a single decade midpoint. In the traditional case, 2040 and 2049 would be considered to  
251 be members of the same class since they are in the same decade, and information would  
252 be lost as there is no way to distinguish whether the samples come from the beginning  
253 or the end of the decade. Using fuzzy encoding, this information of where a sample lies  
254 in each decade is retained. We use a triangular membership function (Zadeh, 1965) with  
255 a width equal to one decade such that each year has partial membership in one or two  
256 neighboring decade classes, and the total membership sums to one. Following this method,  
257 any year directly on a decade midpoint has membership in that class only while years  
258 that fall between decade midpoints have membership in the two neighboring classes. The  
259 year 1925, for example, is mapped to a likelihood of one for the class 1925 and a like-  
260 lihood of zero in all other classes. The year 2078 is mapped to a likelihood of 0.7 for the

261 2075 class and a likelihood of 0.3 for the 2085 class. Note that decoding class likelihoods  
 262 back to their year is simply the decade-weighted sum of the likelihood:  $0.7 \times 2075 + 0.3$   
 263  $\times 2085 = 2078$ . A visualization of the encoding/decoding process can be found in Fig.  
 264 2 of Barnes et al. (2020).

### 265 3.2 Neural Network Training

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

282 The neural networks have several hidden nodes which enable them to learn com-  
 283 plicated relationships between the input and output data. However, with limited train-  
 284 ing data, many of these learned relationships will capture patterns of the noise in the  
 285 training dataset which can lead to overfitting (Srivastava et al., 2014). Atmospheric sci-  
 286 ence data is also highly correlated in space and this collinearity can cause complications  
 287 in the interpretation of the learned weights (Newell & Lee, 1981). Thus, to reduce over-  
 288 fitting and address these issues, we apply ridge regularization ( $L_2$  regularization, see Barnes  
 289 et al., 2020) to the weights of the first hidden layer. Ridge regularization adds a penalty  
 290 (called the ridge penalty) to the square of the weights so the solution is penalized for hav-  
 291 ing large weights. Through training, this acts to shrink the largest weights, thus spread-  
 292 ing the weight out more evenly across multiple grid points. In our application this re-  
 293 sults in a more even distribution of importance across regions with strong spatial cor-  
 294 relation and improves the performance of the neural networks when given data they were

295 not trained on, namely those models in the testing set (elaborated on in Fig. 3, Section  
296 4 of Barnes et al., 2020).

297 Unlike classical approaches which tune the neural network to reduce the mean squared  
298 error (MSE) between the predicted and truth outputs in the testing set (in our case this  
299 would be the MSE between the truth and predicted years), we select the ridge penalty  
300 that minimizes the time of emergence of the forced climate signal (see Section 3.3). Us-  
301 ing time of emergence, rather than MSE, to identify the appropriate ridge penalty en-  
302 sures that we are encouraging the neural networks to learn the patterns of the forced re-  
303 sponse across all decades. When a small ridge penalty is used, the neural networks are  
304 able to predict the year at the end of the 21st century almost perfectly, at the expense  
305 of the predictive skill in earlier decades. This results in a later calculation of time of emer-  
306 gence for the testing set. Slightly increasing the ridge penalty can allow the neural net-  
307 works to detect the climate change signal slightly earlier (Figure S4). The ridge penalty  
308 used for each input vector can be found in the supplementary materials. We use the same  
309 ridge penalty for all 100 neural networks trained on each input vector.

310 All input fields (for climate models and observations) are standardized to assist with  
311 the training and overall performance of the neural network. We subtracted the 1980–2019  
312 mean at each grid point of the input fields for each climate model independently. This  
313 recasts each input field to measure the change relative to the 1980–2019 mean, rather  
314 than the raw magnitudes, which improves the predictive skill of the neural networks and  
315 is also appropriate for identifying indicator patterns of forced change. Since values for  
316 precipitation change are often on the order of  $10^{-6}$ , while the values for temperature change  
317 are on the order of  $10^0$ , we normalized the data so the inputs to the neural network all  
318 have a similar magnitude. To do this, the data from 1980–2019 at each grid point for  
319 each climate model are detrended using ordinary least squares linear regression. We then  
320 take the multi-model mean of the standard deviation of the detrended 1980–2019 data  
321 for each grid point. The input fields are then divided by this new field of standard de-  
322 viations so the inputs are of the same magnitude and fall in a reasonable range for train-  
323 ing the neural networks. Since all our observational datasets include the years 1980 to  
324 2019, we standardize the observations as if they were additional climate models: raw ob-  
325 servations are subtracted by their own 1980–2019 mean, and divided by the same multi-  
326 model standard deviations that were used to standardize the CMIP6 data.

### 327 **3.3 Time of Emergence Calculation**

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

330 climate by the neural network. Specifically, we define the TOE as the year when the neu-  
 331 ral network is able to distinguish that year’s map from any map over a historical base-  
 332 line period. In this work we define this baseline period as 1920–1959 and, under this def-  
 333 inition, the earliest possible TOE estimate is 1960. The TOE is estimated for each cli-  
 334 mate model simulation independently and a schematic of how the TOE is estimated is  
 335 presented in Figure 2. First, we calculate the maximum of the neural network-predicted  
 336 years over 1920–1959 for each model, which is referred to as the baseline maximum. We  
 337 then identify the TOE as the earliest year in which a map, and all subsequent maps, per-  
 338 manently exceed the baseline maximum. In Figure 2, sample model 1 has a baseline max-  
 339 imum of 1966 and permanently exceeds this prediction threshold in 2028. Sample model  
 340 2 has a baseline maximum of 1981 and permanently exceeds this threshold in 1989. Thus,  
 341 the TOE for sample models 1 and 2 are estimated as 2028 and 1989, respectively. In the  
 342 following sections we present the TOE for the testing set, however TOE estimates are  
 343 similar for both the training and testing sets.

### 344 **3.4 Layer-wise Relevance Propagation**

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

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

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

366 There are several LRP decomposition rules which provide different methods of vi-  
 367 sualizing neural networks (Lapuschkin, 2019; Mamalakis et al., 2021). In our applica-  
 368 tions we use the  $\alpha\beta$ -rule which propagates positive relevance (regions that act to increase  
 369 the class likelihood) and negative relevance (regions that act to decrease the class like-  
 370 lihood) separately. Using the parameters  $\alpha = 1$  and  $\beta = 0$  we choose to only propa-  
 371 gate positive relevance, thus highlighting the regions that added to the likelihood of the  
 372 selected decade class. We also looked at the relevance maps for  $\beta = 1$  and found that  
 373 propagating negative relevance did not impact the conclusions.

### 374 **3.5 Signal-to-Noise Ratio Calculation**

375 In Section 4, we compare the LRP relevance maps to maps of signal-to-noise ra-  
 376 tio (S/N ratio), a more conventional method for identifying indicator patterns of the forced  
 377 response. S/N ratio consists of three distinct components: the forced signal, which is di-  
 378 vided by the sum of noise due to internal variability, and noise due to climate model dis-  
 379 agreement. A higher S/N ratio indicates that the signal of the forced response within  
 380 the climate models is very large relative to the underlying noise. We evaluate the S/N  
 381 ratio for each grid point separately, following the methodology in Hawkins and Sutton  
 382 (2012). First, we smooth the data from 1920 to 2098 for each climate model using a fourth-  
 383 order polynomial fit. The signal is defined as the difference between 2090 and 1920 in  
 384 the smoothed data, while internal variability is defined as the standard deviation of the  
 385 residuals from the smoothed data, and climate model disagreement is defined as the stan-  
 386 dard deviation of the signals calculated for all the climate models. S/N ratio is calcu-  
 387 lated by dividing the climate signal by the 90% confidence interval in the noise: inter-  
 388 nal variability and climate model disagreement. S/N ratio, and its components, can be  
 389 seen in Figure S8.

## 390 **4 Global Precipitation and Temperature**

### 391 **4.1 Time of Emergence**

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

397 ral networks are able to detect the patterns of the forced response and distinguish be-  
 398 tween maps in different years. These patterns of the forced response detected by the neu-  
 399 ral networks are generalizable across CMIP6 models, and as a result the neural network  
 400 has predictive skill for seen data (the training set, see the supplementary materials) as  
 401 well as unseen data (the testing set). These behaviors are shown in Figure 3 which presents  
 402 the predicted years from one trained neural network for each combination of global pre-  
 403 cipitation and temperature input fields. Across all input vectors, a similar story of the  
 404 forced signal unfolds. Prior to the TOE, the neural network is unable to identify pat-  
 405 terns that allow it to accurately predict the year. As a result, the neural network is equally  
 406 confident (or unconfident) about which year, between 1920 and the TOE, each input came  
 407 from, so it predicts years right around the middle of the 20th century. After the TOE,  
 408 the predicted years tend to follow a 1:1 line with the truth years, indicating that the neu-  
 409 ral network has identified reliable indicators of change for this period.

410 Although the neural networks are trained on climate model simulations, their learned  
 411 patterns can be used to predict the year for observational data as well. When observa-  
 412 tions are used as input, the predicted years increase with time, just as they do for cli-  
 413 mate model input (Figure 3). This means that the indicators of change derived by the  
 414 neural networks trained on climate models simulations are largely consistent with the  
 415 real world. Pearson correlations ( $r$ ) of the actual years with the years predicted by each  
 416 neural network are shown in Figure 4. All correlations are positive, indicating that the  
 417 years predicted by the neural networks increase with time. These correlations are strongest  
 418 for temperature and combined observations ( $r \approx 0.9$ ), but still quite high for precipita-  
 419 tion ( $r \approx 0.8$ ). Correlations of actual years with predicted years are slightly higher for  
 420 the combined temperature and precipitation observations than for temperature obser-  
 421 vations alone (Figure S5), suggesting that the multivariate indicator patterns derived from  
 422 climate model data are useful for understanding trends in the present-day climate. Across  
 423 all variables, the highest observational correlations are found by the neural networks trained  
 424 on seasonal-mean data. The correlation of actual years with predicted years for precip-  
 425 itation observations are sensitive to the dataset of choice, which is expanded on in Sec-  
 426 tion S4 and Figures S5 and S6.

427 The average TOEs, calculated from the climate models in the testing sets of all 100  
 428 trained neural networks for each input field (Figure 5), reveal that the forced response  
 429 can be detected earlier in maps of temperature than in maps of precipitation (Figure 5a-  
 430 c). When presented with combined fields the neural networks are, in many cases, able  
 431 to detect the forced signal even earlier than when given single fields alone (Figure 5b,f).  
 432 The TOE is generally earlier for the neural networks trained on seasonal-mean data than

433 for the neural networks trained on annual-mean data (Figure 5d-f). This is most notable  
 434 for precipitation fields, likely because there are large seasonal precipitation responses muted  
 435 by taking the annual mean (Tabari & Willems, 2018; Zappa et al., 2015). The TOEs are  
 436 earlier for temperature and precipitation combined than temperature alone when using  
 437 seasonal-mean maps (Figure 5b), but are approximately equal when using annual-mean  
 438 or monthly-mean maps (Figure 5a,c), which suggests that precipitation only improves  
 439 upon the detectability of the forced temperature signal when seasonal-mean fields are  
 440 used. While annual-mean precipitation may mute seasonal precipitation signals, monthly-  
 441 mean precipitation is noisy. In this case, seasonal means emerge as the appropriate tem-  
 442 poral segments for detecting precipitation change, underlining the importance for the  
 443 intentional and intelligent selection of neural network inputs.

444 The neural networks identify the earliest TOEs when trained on seasonal-mean tem-  
 445 perature and precipitation combined (Figure 5b,f). The TOE results for all 100 seasonal-  
 446 mean neural networks are summarized in the box plots in Figure S7. While the improve-  
 447 ment in forced response detection is small when precipitation is combined with temper-  
 448 ature, it is still notable given that the forced signal of temperature is much clearer than  
 449 the forced signal of precipitation. We use these variables as an initial example for em-  
 450 ploying this neural network methodology. We anticipate that more robust results might  
 451 be found for combinations of variables that have more distinct combined signals, such  
 452 as humidity and temperature (Fischer & Knutti, 2012).

## 453 **4.2 Indicator Patterns for Combined Variables**

454 Having shown that the neural networks are able to predict the year given seasonal  
 455 means of temperature and precipitation (Figures 3, 5), we now identify and explore the  
 456 spatial indicator patterns used by the neural networks to make correct predictions. By  
 457 understanding the neural networks' decision-making process, we can identify which re-  
 458 gions act as combined (multi-seasonal and multi-variable) indicators of forced change amidst  
 459 a background of internal variability and climate model disagreement. To identify these  
 460 indicator patterns, we apply LRP to all climate model samples in the training and test-  
 461 ing sets from the year 2090 for the seasonal-mean combined neural networks. Averag-  
 462 ing the LRP results for each season and variable, we highlight the regions that have the  
 463 highest mean relevance across the 37 CMIP6 climate models and 100 trained neural net-  
 464 works. The relevance maps for temperature (precipitation) are shown in Figure 6a-d (7a-  
 465 d) and indicate the importance of each region in the neural networks' predictions of the  
 466 year 2090.

467 LRP identifies temperature over North Africa and Central Asia in JJA (Figure 6c)  
 468 and the Andes and Central Africa in SON (Figure 6d) as the most relevant regions for  
 469 predicting the year. For precipitation, the regions of highest relevance can be found in  
 470 Canada and Russia in DJF and SON (Figure 7a,d) and in Central Africa and India in  
 471 JJA and SON (Figure 7c,d). That is to say that these are the regional patterns iden-  
 472 tified by the neural networks that indicate the presence of forced change across the CMIP6  
 473 climate models. The scale of the color bars are different between Figures 6 and 7, such  
 474 that the darkest regions in the temperature maps are approximately one order of mag-  
 475 nitude more relevant than the darkest regions in the precipitation maps. Hence, the neu-  
 476 ral network is relying more heavily on the temperature inputs than the precipitation in-  
 477 puts in order to accurately predict the year. This is not surprising because the forced  
 478 signal of temperature is clearer than the forced signal of precipitation (Fig. SPM.7 in  
 479 Field et al., 2014). Even so, including seasonal precipitation allows the neural networks  
 480 to detect forced change earlier within combined fields than in temperature fields alone  
 481 (Figure 5b). The improvement in neural network performance provided by precipitation  
 482 (alongside temperature) is particularly noteworthy given that the S/N ratio for temper-  
 483 ature is larger than the S/N ratio for precipitation in all seasons and regions (Figures  
 484 6e-h, 7e-h, discussed further in this section). In other words, the forced temperature sig-  
 485 nal is always more pronounced than the forced precipitation signal, but the precipita-  
 486 tion signal is still useful for detecting forced change.

487 LRP is designed to highlight the regions that were most relevant for predicting the  
 488 correct class (in our case, the correct decade class). These LRP indicator patterns for  
 489 2090 are not the time-mean patterns of the forced response, they are the patterns used  
 490 by the neural network to distinguish the end of the 21st century from all other decades.  
 491 This is distinctly different from S/N ratio which identifies the regions where the forced  
 492 change from 1920 to 2090 is largest relative to internal variability and climate model spread.  
 493 Maps of S/N ratio for temperature are shown in Figure 6e-h, and the corresponding maps  
 494 for precipitation are shown in Figure 7e-h, where a higher S/N ratio (darker green) in-  
 495 dicates a clearer forced signal. These regions of high S/N ratio are consistent with other  
 496 related studies (e.g., Hawkins et al., 2020). For the most part, the indicator patterns iden-  
 497 tified by LRP correspond with the regions with the highest S/N ratios. Calculating the  
 498 Spearman’s rank correlation ( $\rho$ ) between each map of relevance and S/N ratio, we find  
 499 that there is generally a strong positive correlation ( $0.71 \leq \rho \leq 0.77$ ) between the LRP  
 500 indicator patterns and the S/N ratios for temperature, and a moderate positive corre-  
 501 lation ( $0.30 \leq \rho \leq 0.56$ ) for precipitation. The exact correlation coefficients between  
 502 each map are displayed in the subtitles for Figures 6e-h and 7e-h.

503           Given that precipitation only contributes a small amount of relevance compared  
504 to temperature, it is perhaps unsurprising that there are several regions where the S/N  
505 ratio for precipitation is high, but the relevance is low (e.g. Alaska in JJA, Figure 7c,g  
506 or South Africa in SON, Figure 7d,h). Most likely, the forced signal of temperature is  
507 clear enough that these regions do not add to the predictive skill of the neural networks.  
508 Regions also exist where the S/N ratio for temperature is high despite low relevance (e.g.  
509 North Africa in DJF, Figure 6a,e), although these are more rare, as hinted by the strong  
510 correlation between the temperature maps of S/N ratio and relevance. In contrast, there  
511 are fewer regions with high relevance despite low S/N ratios, but they do occur (e.g. SON  
512 temperatures in northern South America, Figure 6d,h). These high-relevance, low-S/N  
513 ratio regions confirm that the indicator patterns identified by LRP capture more than  
514 the local S/N ratio. Some reasons a region/variable/season may be important in terms  
515 of LRP, but not in terms of S/N ratio, are: 1) LRP may be identifying places in our data  
516 where a signal exists only in the combination of regions/seasons/variables, which would  
517 not be captured by this definition of S/N ratio. 2) Since LRP highlights the patterns the  
518 neural networks use to predict the correct decade over all other decades, it may be cap-  
519 turing abrupt nonlinear changes in the climate that are filtered out by the century-long  
520 analysis of S/N ratio. In the next section, we discuss further applications of neural network-  
521 derived indicator patterns and task the network with the much harder problem of iden-  
522 tifying changes in extreme precipitation over the Americas.

## 523   **5 Extreme Precipitation over the Americas**

524           We now task the neural networks to predict the year given combinations of sea-  
525 sons for a single variable: extreme precipitation over the Americas. We choose to shift  
526 our focus for a few reasons. First, we wish to demonstrate that this neural network ap-  
527 proach can be extended to variables that have considerable noise (like extreme precip-  
528 itation, see Figure S8), and datasets that do not cover the globe. Second, extreme pre-  
529 cipitation has major implications for human health (Ali et al., 2019; Eekhout et al., 2018;  
530 Rosenzweig et al., 2002) but there is considerable disagreement between climate mod-  
531 els in its signal (Figure S8). This neural network approach can be used to identify agreed-  
532 upon patterns despite climate model spread. Further in this section, we will demonstrate  
533 that LRP maps can be used to investigate climate model differences and better under-  
534 stand the time evolution of the forced response.

535           The extreme precipitation signal is not as pronounced as the temperature signal,  
536 and using the Americas rather than the full globe limits the amount of unique informa-  
537 tion in the input field. Nevertheless, the neural networks are still able to detect patterns

538 of forced change. Figure 8 depicts the years predicted by one neural network trained on  
539 seasonal-mean extreme precipitation. As in Figure 3, the neural network is unable to ac-  
540 curately predict the year given CMIP6 data prior to the TOE around 2010, whereafter  
541 the predicted years generally follow the 1:1 line with the truth years, indicating that the  
542 neural network has identified reliable indicators of change for this period. All Pearson  
543 correlations of the actual years with the predicted years for extreme precipitation in ob-  
544 servations are positive ( $r \approx 0.4$ ), demonstrating that the indicator patterns found in cli-  
545 mate models can be successfully applied to observations (Figure 4). These correlations  
546 are not as strong as those for mean precipitation observations, due in part to the mag-  
547 nitude of climate model disagreement in extreme precipitation as well as the observa-  
548 tional dataset used: ERA5. As shown in Figure S6, the correlations of actual with pre-  
549 dicted years for ERA5 precipitation observations are far smaller than those for GPCP  
550 observations. ERA5 tends to perform poorly in remote regions such as northern North  
551 America and northwestern South America (Bell et al., 2021), which may be responsible  
552 for these low correlations. The correlation between actual years and neural network-predicted  
553 years for extreme precipitation observations are explored in much more detail by Madakumbura  
554 et al. (2021).

555 To investigate the indicator patterns used by the neural networks to predict the  
556 year when the forced signal first emerges from the background noise, we apply LRP to  
557 all climate model samples in the training and testing sets for all 100 neural networks at  
558 the TOE (using the TOE calculated for each climate model and neural network individ-  
559 ually, see Figure S9). LRP points to western South America in DJF and British Columbia  
560 in MAM and SON as the most relevant regions when the neural networks first detect the  
561 forced response (Figure 9a-d). These LRP maps exhibit a more even distribution in rel-  
562 evance across each region and season than the end-of-the-21st-century LRP maps of global  
563 temperature and precipitation (Figures 6a-d, 7a-d). Predicting the year at the TOE, when  
564 the signal has just barely emerged from the background climate, likely requires the neu-  
565 ral networks to use all of the information available to them.

566 Up to this point, we have only considered the mean LRP maps across climate mod-  
567 els. Since the neural networks are nonlinear by nature, they can identify multiple pat-  
568 terns that differ between climate models for a given decade. We apply k-means cluster-  
569 ing to all 3200 LRP maps at the TOE (32 climate models samples, 100 neural networks)  
570 to identify two distinct indicator patterns that are being used by the climate models (Fig-  
571 ure 9e-l, see the supplementary materials for more details on k-means clustering). Tak-  
572 ing the difference between the mean LRP maps for clusters one and two reveals that the  
573 Amazon in JJA is a highly relevant region in cluster one, while western Canada in DJF

574 is a highly relevant region in cluster two (Figure 9m-p). With the sole exception of MPI-  
575 ESM1-2-HR, all 100 LRP maps for each individual climate model fall cleanly into one  
576 cluster or the other, suggesting that there are two distinct ways in which the forced sig-  
577 nal emerges in the CMIP6 simulations (Figure 10). Interestingly, when k-means is in-  
578 structed to identify 32 unique clusters within the LRP maps, each cluster contains all  
579 100 relevance maps for each of the 32 climate models. In other words, the pathway used  
580 by the neural networks to predict the year is unique to each climate model and distin-  
581 guishable from all other climate models, regardless of whether the climate model sam-  
582 ples appear in the training or testing sets (further investigated by Labe & Barnes, 2022).

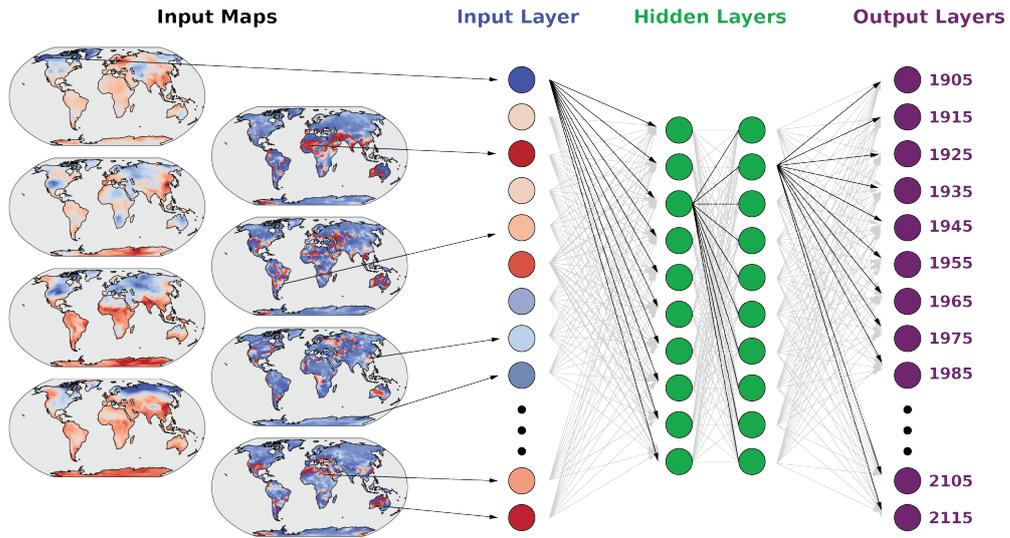
583 In the same way that indicator patterns can differ between models, indicator pat-  
584 terns are also able to evolve through time (e.g., Barnes et al., 2020; Labe & Barnes, 2021;  
585 Madakumbura et al., 2021). Comparing the LRP maps at the TOE (Figure 11a-d) with  
586 those at the end of the 21st century (Figure 11e-h) highlights the regions that become  
587 more important for predicting the year over time. The difference plots in Figure 11i-l  
588 reveal that the neural network learns to focus on Alaska during MAM, JJA, and SON,  
589 Greenland in JJA and SON, and Quebec in MAM and SON as the forced response be-  
590 comes stronger. These regions are more important for predicting the year at the end of  
591 the 21st century than the early 21st century. While further exploration is required, there  
592 are several reasons a region may become more relevant over time. For example, it may  
593 be that the region does not initially have a clear forced signal, but following some abrupt  
594 change (e.g. an ice-free Arctic) the forced signal becomes extremely pronounced. It may  
595 also be that the region has a signal that is consistently agreed upon by the majority of  
596 CMIP6 climate models, and becomes more relevant compared to other regions as climate  
597 model projections in those other regions drift apart. These time-varying patterns sup-  
598 port the idea that combined indicators are effective for identifying dynamically evolv-  
599 ing patterns of forced change.

## 600 **6 Conclusions**

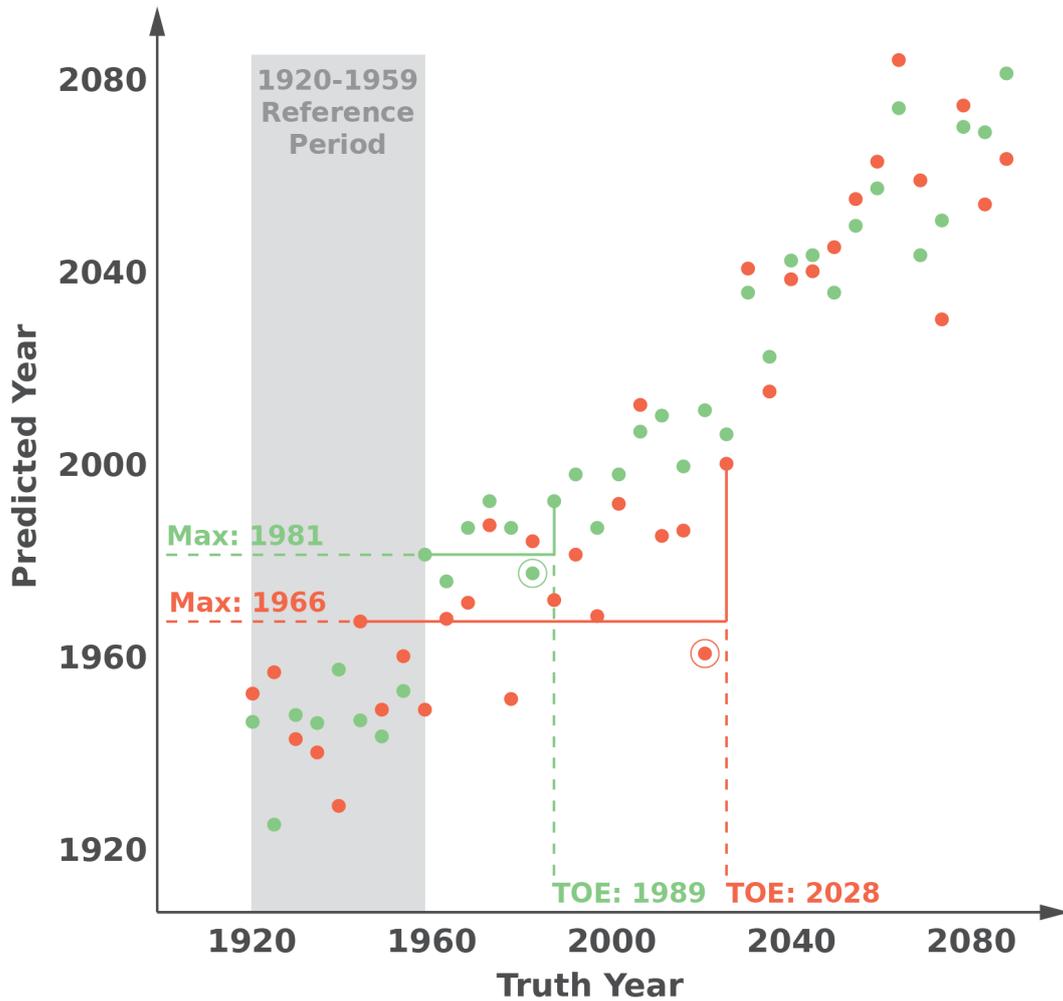
601 Neural networks are powerful tools for identifying patterns of forced change in the  
602 climate system. When tasked with predicting the year given climate model simulations  
603 of temperature, precipitation, or extreme precipitation, artificial neural networks can learn  
604 these patterns of forced change that allow them to distinguish between maps from dif-  
605 ferent years. In combined fields, such as multiple variables, seasons, or both, the forced  
606 response can be detected earlier than in single fields alone. By visualizing the decision-  
607 making process of the neural networks with an explainability method we extracted re-  
608 liable, multivariate patterns of forced change. These neural network-derived combined

609 indicator patterns are complex and nonlinear and capture more than the local signal-  
610 to-noise ratio. Explainability methods take a huge step towards disentangling the rela-  
611 tionships learned by neural networks by pointing out what inputs contributed most to  
612 the final prediction, but they stop short of explaining why.

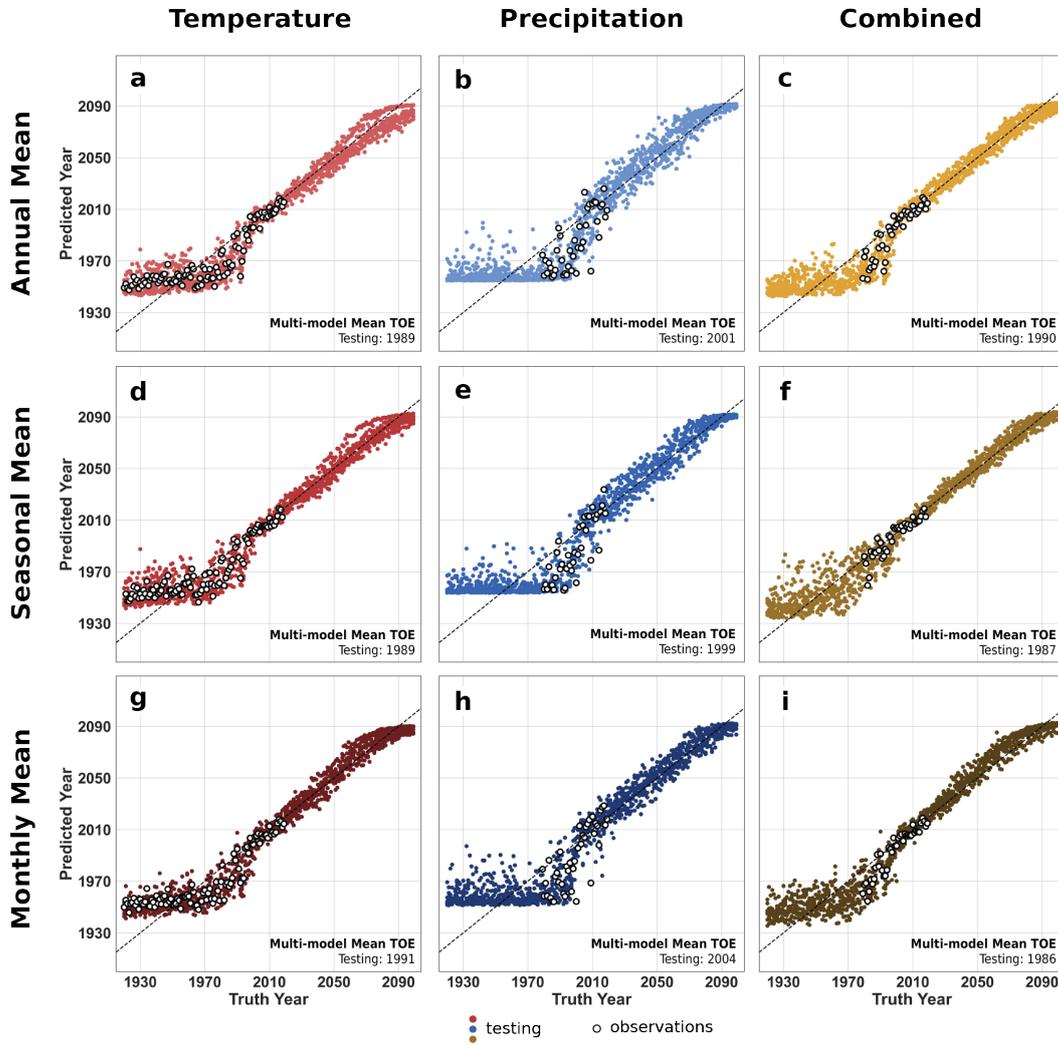
613 Expanding on previous work by Barnes et al. (2020), we used k-means clustering  
614 in tandem with layer-wise relevance propagation to study the relationships learned by  
615 the neural networks. This approach revealed two distinct ways in which the extreme pre-  
616 cipitation response emerges in CMIP6 data. While combining neural network explain-  
617 ability methods with other statistical techniques can enhance explanations of neural net-  
618 work decisions, there is still a large gap between what the neural network has learned  
619 and what we can explain post hoc. Some unanswered questions, such as “why does tem-  
620 perature in Region A combine with precipitation in Region B to improve the signal of  
621 the forced response?” may be better answered with a different architectural approach,  
622 such as neural network designs that are inherently interpretable and do not require post-  
623 hoc approaches like LRP (Rudin, 2019). This is a natural next step for future work. The  
624 flexibility and accessibility of this framework provide several other future research di-  
625 rections. Given that this predict-the-year approach can be applied to observational data,  
626 one possible extension of this work could involve exploring the observed features of forced  
627 change that are consistent with climate model simulations. There is also space for these  
628 methods to be used to determine which definitions of seasons are optimal for detecting  
629 forced change. While we used meteorological seasons here, there may be more appropri-  
630 ate definitions, such as unique definitions of the wet and dry seasons, or the shoulder sea-  
631 sons, that vary between variables and regions. Furthermore, this framework should be  
632 expanded to other variables, regions of focus, and climate change scenarios, to identify  
633 the combined indicators that best elucidate the forced signal. For example, extreme pre-  
634 cipitation and extreme drought may combine to capture the increased volatility in pre-  
635 cipitation extremes that are expected with climate change (O’Gorman, 2015). Further  
636 application of this technique to compound climate extremes, such as heat wave inten-  
637 sity, drought duration, and flood frequency, may reveal that explainable neural networks  
638 are useful for assessing societal impacts and improving climate change preparedness.



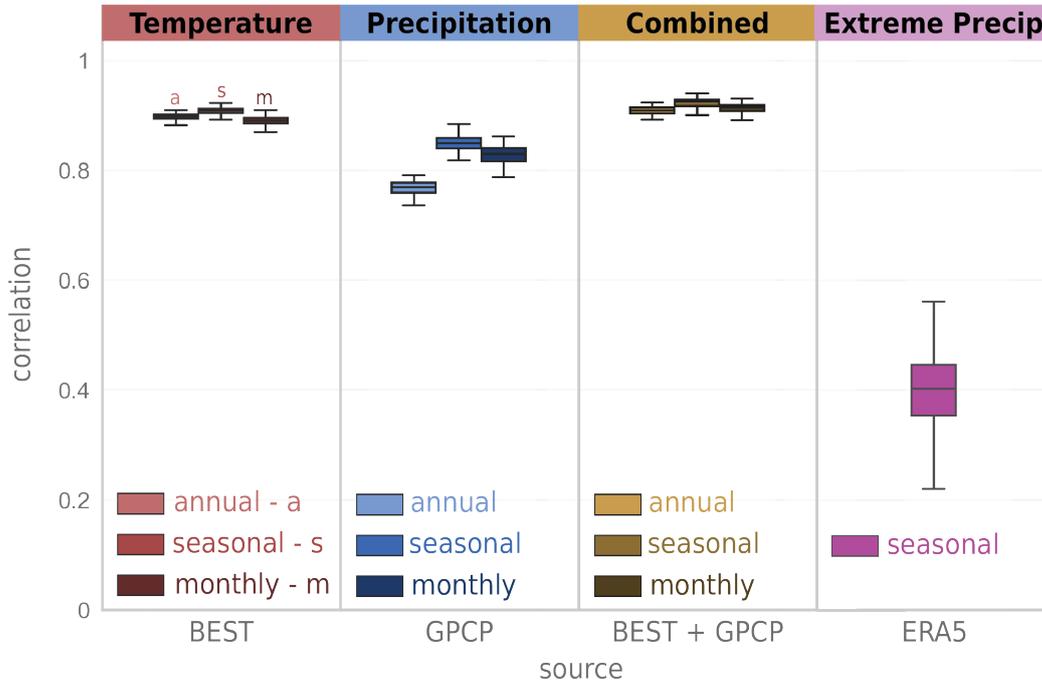
**Figure 1. Schematic of the fully connected neural network architecture.** Inputs from multiple maps of data are flattened into an input layer vector (size of the input layer ranges from 948 to 22,752). These inputs are fed through two hidden layers with ten nodes each. The neural network is trained 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.



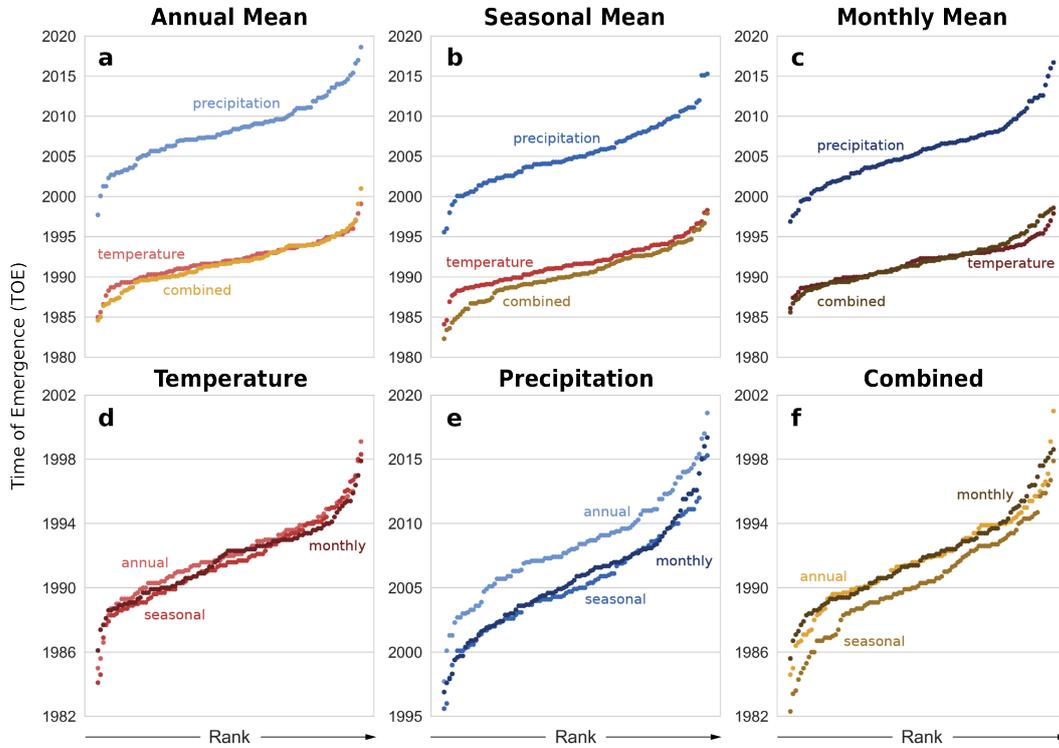
**Figure 2. Calculation of TOE.** The TOE is defined as the earliest year in which a map, and all subsequent maps, permanently exceed the maximum predicted year from the baseline period (1920-1959). The baseline maximum for each model is indicated by the horizontal lines, the last year that falls below 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.



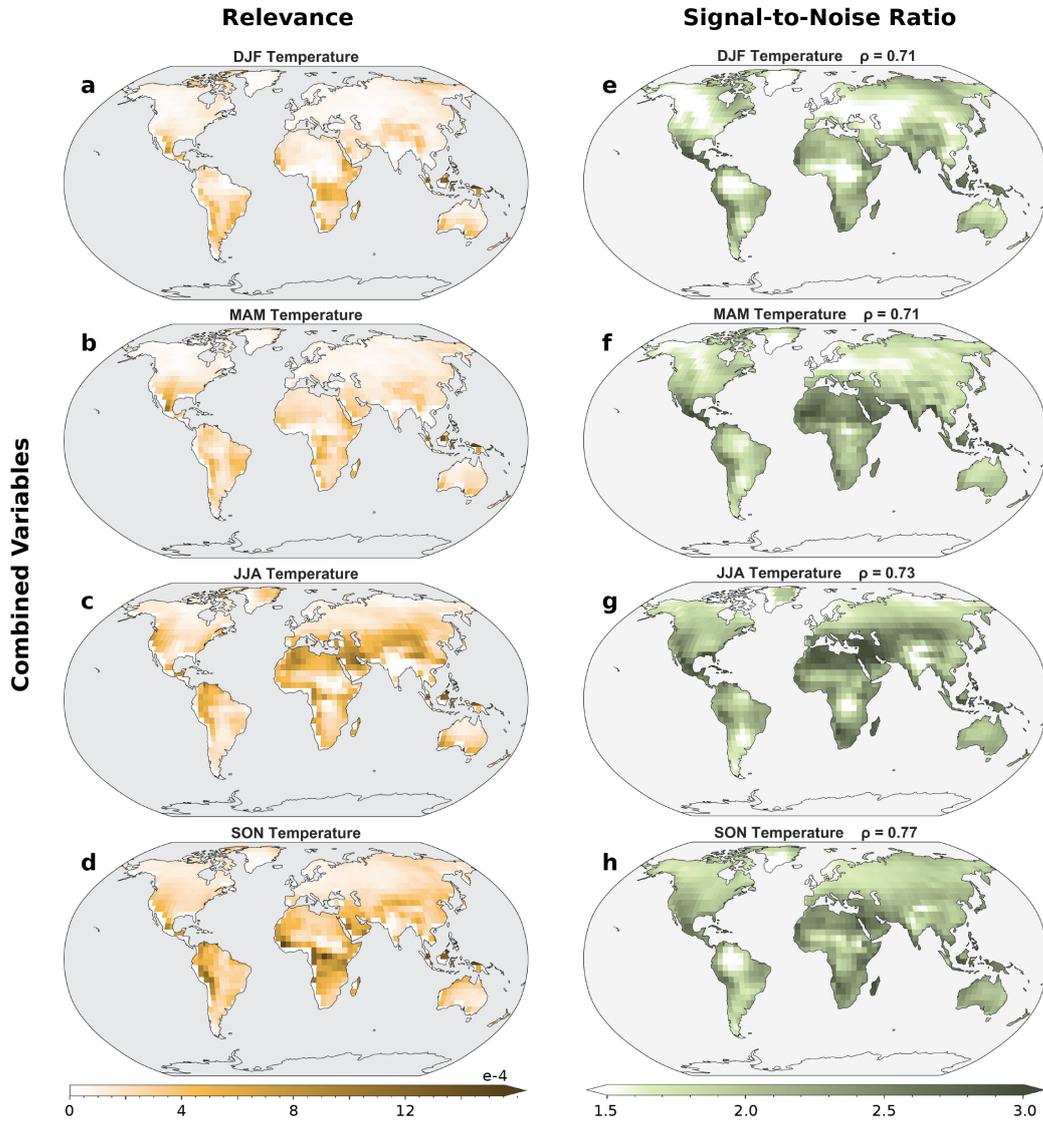
**Figure 3.** Neural network output for temperature and precipitation. Year predicted by the neural network (y-axis) versus the truth year (x-axis) for temperature (a, d, g), precipitation (b, e, h), and temperature and precipitation combined (c, f, i). Input maps include annual-mean data (a, b, c),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



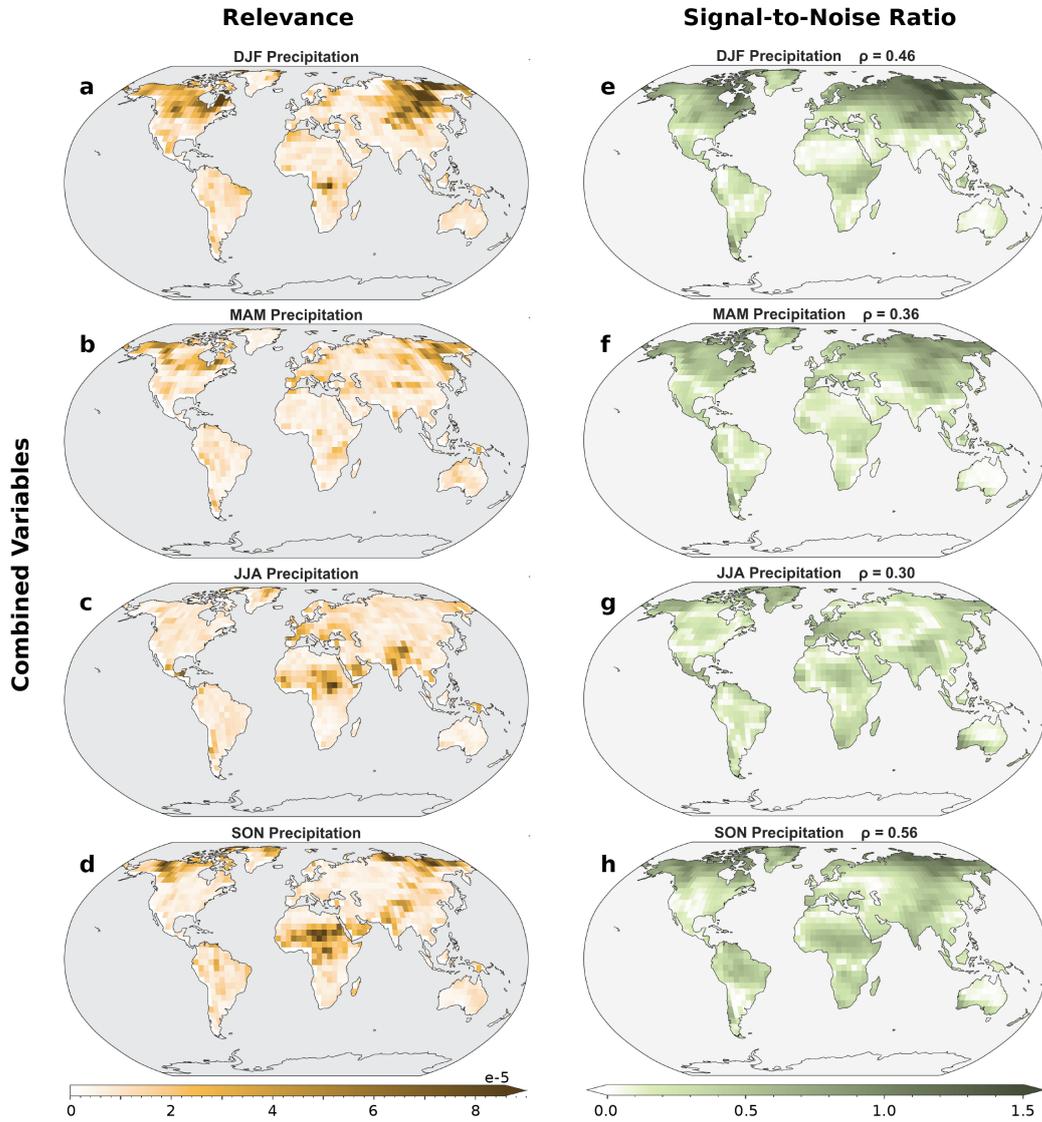
**Figure 4. Correlation of actual years with predicted years for observations.** Pearson correlations of the actual years with the years predicted by 100 trained neural networks given observations of temperature, precipitation, and extreme precipitation. Correlations were computed for all years beginning 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 S5 and S6.



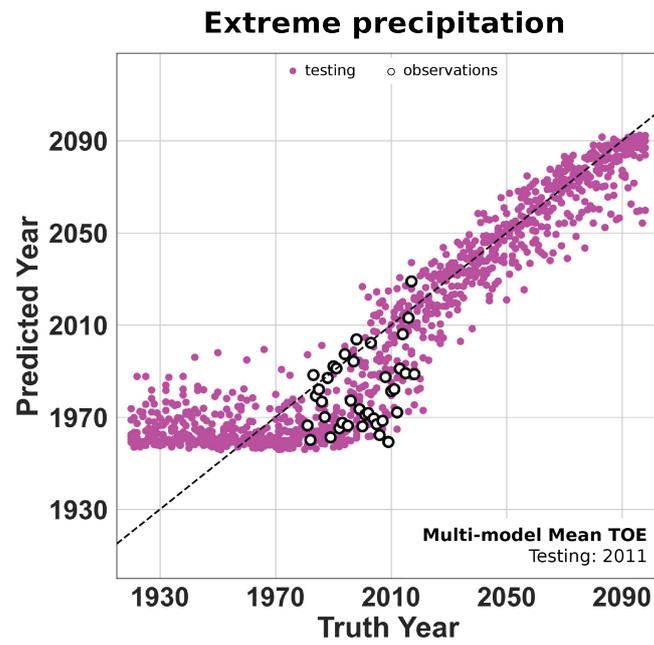
**Figure 5.** Mean TOE for each input field. Comparison of the mean time of emergence identified by neural networks trained on annual-mean (a), seasonal-mean (b), and monthly-mean (c) input fields, and neural networks trained on temperature (d), precipitation (e), and temperature and precipitation combined (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, and that the TOE results for each set of neural networks appear once in the panels a, b, and c, and once in d, e, and f.



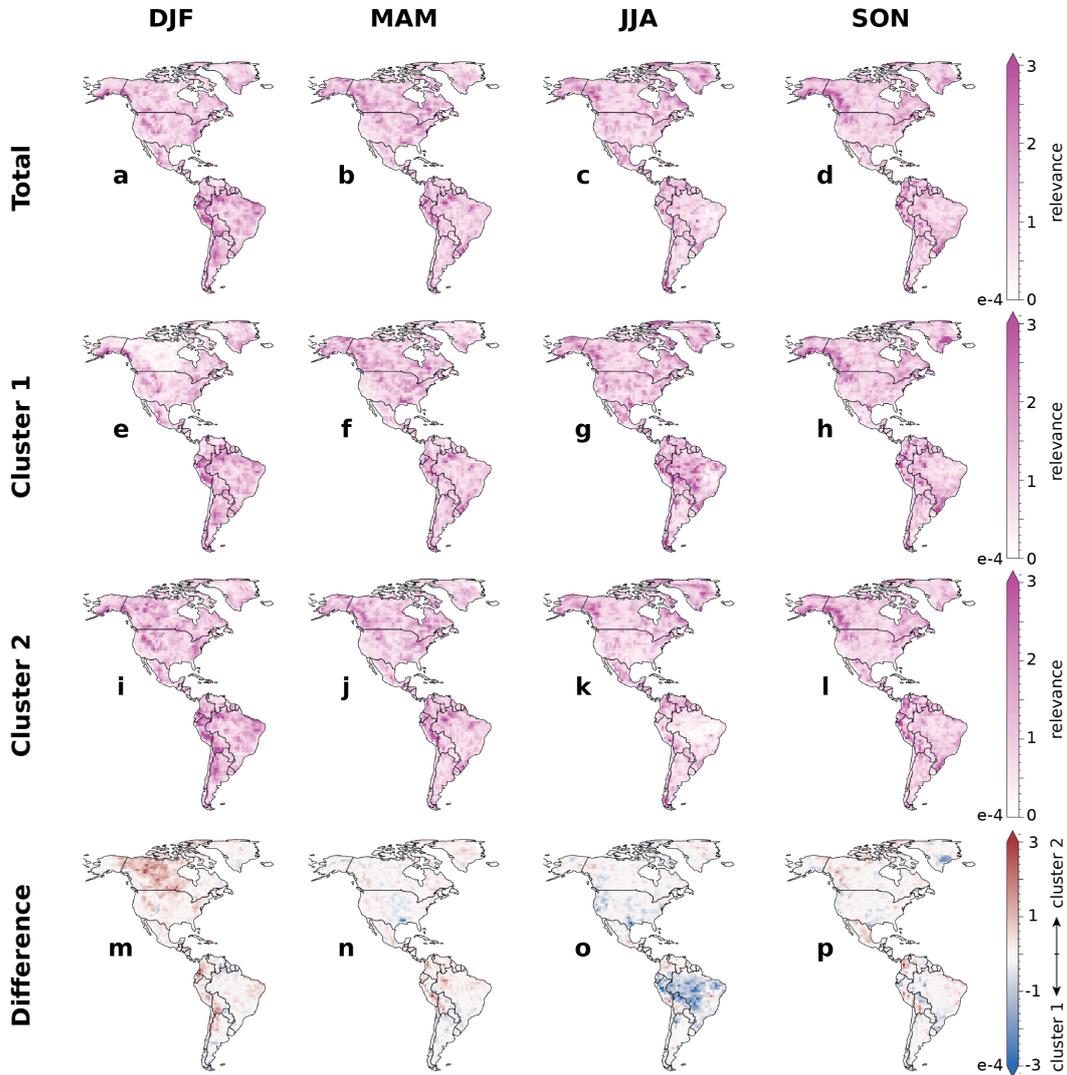
**Figure 6.** Combined indicator patterns of the forced response (temperature). Average temperature LRP results for the seasonal-mean combined neural networks (left, in yellow) and S/N ratio (right, in green) for 2090. Darker shading indicates regions of temperature that are more relevant for the neural network’s prediction or have a higher S/N ratio. The Spearman’s rank correlation ( $\rho$ ) between corresponding maps of relevance and S/N ratio are shown in the subtitles of panels e-h.



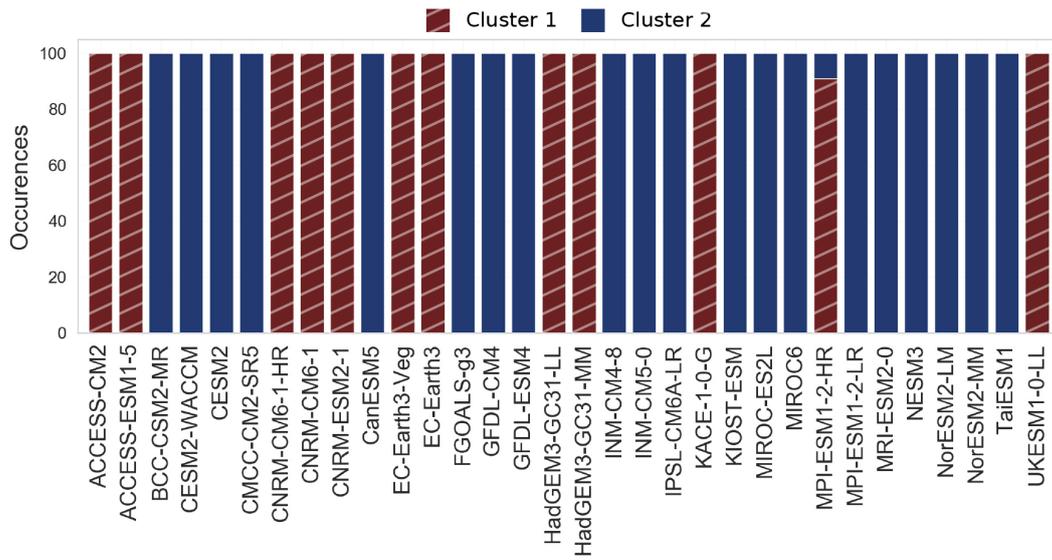
**Figure 7.** Combined indicator patterns of the forced response (precipitation). Average precipitation LRP results for the seasonal-mean combined neural networks (left, in yellow) and S/N ratio (right, in green) for 2090. Darker shading indicates regions of precipitation that are more relevant for the neural network’s prediction or have a higher S/N ratio. The Spearman’s rank correlation ( $\rho$ ) between corresponding maps of relevance and S/N ratio are shown in the subtitles of panels e-h.



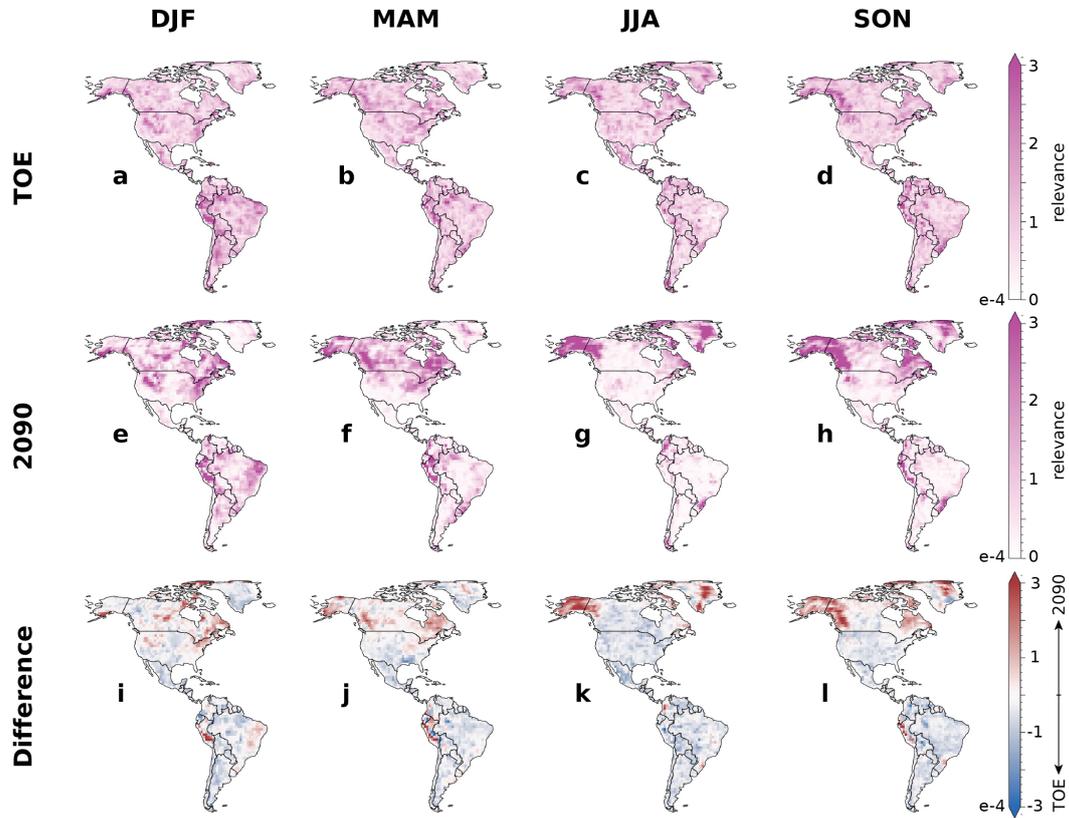
**Figure 8.** Neural network output for extreme precipitation. Year predicted by the neural network (y-axis) versus the truth year (x-axis) given seasonal-mean maps of extreme precipitation. Testing data is shown in pink and observations are shown in white.



**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. Note that panels a-d are identical to panels a-d in Figure 11.



**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. Note that panels a-d are identical to panels a-d in Figure 9.

## 639 **Acknowledgments**

640 We thank the two anonymous reviewers, the Editor, Maria Rugenstein, and Jessica Witt  
 641 for their constructive feedback, which substantially improved our study. I would also like  
 642 to thank the members of the Barnes Group for all the science discussions that brought  
 643 about new ideas, and my parents and siblings for their support through the pandemic.  
 644 This material is based upon work supported by the U.S. Department of Energy, Office  
 645 of Science, Office of Advanced Scientific Computing Research, Department of Energy Com-  
 646 putational Science Graduate Fellowship under Award Number DE-SC0020347, and by  
 647 NOAA MAPP grant NA19OAR4310289. We acknowledge the World Climate Research  
 648 Programme’s Working Group on Coupled Modelling, which is responsible for CMIP, and  
 649 we thank the climate modeling groups for producing and making available their model  
 650 output. For CMIP the U.S. Department of Energy’s Program for Climate Model Diag-  
 651 nosis and Intercomparison provides coordinating support and led development of soft-  
 652 ware infrastructure in partnership with the Global Organization for Earth System Sci-  
 653 ence Portals.

## 654 **Data Availability Statement**

655 All data used in this study is publicly available and referenced throughout the paper.  
 656 The CMIP6 simulations used in this study can be via the Earth System Grid Federa-  
 657 tion (ESGF, <https://esgf-node.llnl.gov/projects/cmip6/>). Monthly temperature obser-  
 658 vations are available through Berkeley Earth (<http://berkeleyearth.org/data/>). Global  
 659 Precipitation Climatology Project monthly global precipitation fields are available through  
 660 the NOAA Physical Sciences Laboratory (<https://psl.noaa.gov/data/gridded/data.gpcp.html>).  
 661 Monthly, daily, and sub-daily precipitation reanalyses were provided by the European Cen-  
 662 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)  
 663 [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)  
 664 [data/jra-55](https://climatedataguide.ucar.edu/climate-data/jra-55)). Python code used in this work has been made publicly available at <https://github.com/jaminrader/>

## 665 **References**

- 666 Abiodun, O. I., Jantan, A., Omolara, A. E., Dada, K. V., Mohamed, N. A., & Ar-  
 667 shad, H. (2018, November). State-of-the-art in artificial neural network appli-  
 668 cations: A survey. *Heliyon*, 4(11), e00938. doi: 10.1016/j.heliyon.2018.e00938
- 669 Adler, R. F., Sapiano, M., Huffman, G. J., Wang, J., Gu, G., Bolvin, D., . . . Shin,  
 670 D.-B. (2018, April). The global precipitation climatology project (GPCP)  
 671 monthly analysis (new version 2.3) and a review of 2017 global precipitation.  
 672 *Atmosphere*, 9(4). doi: 10.3390/atmos9040138

- 673 Ali, H., Modi, P., & Mishra, V. (2019, September). Increased flood risk in indian  
674 sub-continent under the warming climate. *Weather and Climate Extremes*, *25*,  
675 100212. doi: 10.1016/j.wace.2019.100212
- 676 Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.-R., & Samek, W.  
677 (2015, July). On Pixel-Wise explanations for Non-Linear classifier decisions  
678 by Layer-Wise relevance propagation. *PLoS One*, *10*(7), e0130140. doi:  
679 10.1371/journal.pone.0130140
- 680 Barnes, E. A., Hurrell, J. W., Ebert-Uphoff, I., Anderson, C., & Anderson, D. (2019,  
681 November). Viewing forced climate patterns through an AI lens. *Geophys. Res.  
682 Lett.*, *46*(22), 13389–13398. doi: 10.1029/2019gl084944
- 683 Barnes, E. A., Toms, B., Hurrell, J. W., Ebert-Uphoff, I., Anderson, C., & Ander-  
684 son, D. (2020, September). Indicator patterns of forced change learned by an  
685 artificial neural network. *J. Adv. Model. Earth Syst.*, *12*(9), e2020MS002195.  
686 doi: 10.1029/2020ms002195
- 687 Barnett, T. P., Pierce, D. W., Hidalgo, H. G., Bonfils, C., Santer, B. D., Das, T.,  
688 ... Dettinger, M. D. (2008, February). Human-induced changes in the hy-  
689 drology of the western united states. *Science*, *319*(5866), 1080–1083. doi:  
690 10.1126/science.1152538
- 691 Bell, B., Hersbach, H., Simmons, A., Berrisford, P., Dahlgren, P., Horányi, A., ...  
692 Thépaut, J. (2021). The ERA5 global reanalysis: Preliminary extension to  
693 1950. *Quart J Roy Meteor Soc.*
- 694 Bindoff, N. L., Stott, P. A., AchutaRao, K. M., Allen, M. R., Gillett, N., Gutzler,  
695 D., ... Zhang, X. (2013). Chapter 10 - detection and attribution of climate  
696 change: From global to regional. In *Climate change 2013: The physical sci-  
697 ence basis. IPCC working group I contribution to AR5*. Cambridge, United  
698 Kingdom: Cambridge University Press.
- 699 Bonfils, C. J. W., Santer, B. D., Fyfe, J. C., Marvel, K., Phillips, T. J., & Zimmer-  
700 man, S. R. H. (2020, July). Human influence on joint changes in temperature,  
701 rainfall and continental aridity. *Nat. Clim. Chang.*, *10*(8), 726–731. doi:  
702 10.1038/s41558-020-0821-1
- 703 Brenowitz, N. D., & Bretherton, C. S. (2018, June). Prognostic validation of a  
704 neural network unified physics parameterization. *Geophys. Res. Lett.*, *45*(12),  
705 6289–6298. doi: 10.1029/2018gl078510
- 706 Chen, C., Zhang, P., Zhang, H., Dai, J., Yi, Y., Zhang, H., & Zhang, Y. (2020,  
707 March). Deep learning on Computational-Resource-Limited platforms: A  
708 survey. *Mobile Information Systems*, *2020*. doi: 10.1155/2020/8454327
- 709 Cui, L., Wang, L., Qu, S., Singh, R. P., Lai, Z., & Yao, R. (2019, April). Spa-

- 710 tiotemporal extremes of temperature and precipitation during 1960–2015 in  
 711 the yangtze river basin (china) and impacts on vegetation dynamics. *Theor.*  
 712 *Appl. Climatol.*, *136*(1), 675–692. doi: 10.1007/s00704-018-2519-0
- 713 Dai, A., Lin, X., & Hsu, K.-L. (2007, October). The frequency, intensity, and di-  
 714 urnal cycle of precipitation in surface and satellite observations over low- and  
 715 mid-latitudes. *Clim. Dyn.*, *29*(7-8), 727–744. doi: 10.1007/s00382-007-0260-y
- 716 Donat, M. G., Alexander, L. V., Herold, N., & Dittus, A. J. (2016, October). Tem-  
 717 perature and precipitation extremes in century-long gridded observations,  
 718 reanalyses, and atmospheric model simulations. *J. Geophys. Res.*, *121*(19),  
 719 11,174–11,189. doi: 10.1002/2016jd025480
- 720 Eekhout, J. P. C., Hunink, J. E., Terink, W., & de Vente, J. (2018, April).  
 721 Why increased extreme precipitation under climate change negatively af-  
 722 fects water security. *Hydrol. Earth Syst. Sci. Discuss.*, *22*(11), 1–16. doi:  
 723 10.5194/hess-2018-161
- 724 Eyring, V., Bony, S., Meehl, G. A., Senior, C. A., Stevens, B., Stouffer, R. J., &  
 725 Taylor, K. E. (2016, May). Overview of the coupled model intercomparison  
 726 project phase 6 (CMIP6) experimental design and organization. *Geosci. Model*  
 727 *Dev.*, *9*(5), 1937–1958. doi: 10.5194/gmd-9-1937-2016
- 728 Field, C. B., Barros, V. R., Mastrandrea, M. D., Mach, K. J., Abdrabo, M. A.-K.,  
 729 Adger, N., . . . Yohe, G. W. (2014). Summary for policymakers. In *Climate*  
 730 *change 2014: Impacts, adaptation, and vulnerability. part a: Global and sec-*  
 731 *toral aspects. contribution of working group II to the fifth assessment report of*  
 732 *the intergovernmental panel on climate change* (pp. 1–32). Cambridge, United  
 733 Kingdom and New York, NY, USA: Cambridge University Press.
- 734 Fischer, E. M., & Knutti, R. (2012, September). Robust projections of combined hu-  
 735 midity and temperature extremes. *Nat. Clim. Chang.*, *3*(2), 126–130. doi: 10  
 736 .1038/nclimate1682
- 737 Gaetani, M., Janicot, S., Vrac, M., Famien, A. M., & Sultan, B. (2020, May). Ro-  
 738 bust assessment of the time of emergence of precipitation change in west africa.  
 739 *Sci. Rep.*, *10*(1), 7670. doi: 10.1038/s41598-020-63782-2
- 740 Gettelman, A., Gagne, D. J., Chen, C.-C., Christensen, M. W., Lebo, Z. J., Mor-  
 741 rison, H., & Gantos, G. (2021, February). Machine learning the warm  
 742 rain process. *J. Adv. Model. Earth Syst.*, *13*(2), e2020MS002268. doi:  
 743 10.1029/2020ms002268
- 744 Hawkins, E., Frame, D., Harrington, L., Joshi, M., King, A., Rojas, M., & Sutton,  
 745 R. (2020, March). Observed emergence of the climate change signal: From  
 746 the familiar to the unknown. *Geophys. Res. Lett.*, *47*(6), e2019GL086259. doi:

- 747 10.1029/2019gl086259
- 748 Hawkins, E., & Sutton, R. (2009, August). The potential to narrow uncertainty in  
749 regional climate predictions. *Bull. Am. Meteorol. Soc.*, *90*(8), 1095–1108. doi:  
750 10.1175/2009BAMS2607.1
- 751 Hawkins, E., & Sutton, R. (2012, January). Time of emergence of climate signals.  
752 *Geophys. Res. Lett.*, *39*(1). doi: 10.1029/2011gl050087
- 753 Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J.,  
754 ... Jean-Noël Thépaut (2020, July). The ERA5 global reanalysis. *Quart. J.*  
755 *Roy. Meteor. Soc.*, *146*(730), 1999–2049. doi: 10.1002/qj.3803
- 756 Kim, Y.-H., Min, S.-K., Zhang, X., Sillmann, J., & Sandstad, M. (2020, September).  
757 Evaluation of the CMIP6 multi-model ensemble for climate extreme indices.  
758 *Weather and Climate Extremes*, *29*, 100269. doi: 10.1016/j.wace.2020.100269
- 759 Labe, Z. M., & Barnes, E. A. (2021, June). Detecting climate signals using explain-  
760 able AI with single-forcing large ensembles. *J. Adv. Model. Earth Syst.*, *13*(6),  
761 e2021MS002464. doi: 10.1029/2021ms002464
- 762 Labe, Z. M., & Barnes, E. A. (2022). *Comparison of climate model large ensem-*  
763 *bles with observations in the arctic using simple neural networks.* doi: 10.1002/  
764 essoar.10510977.1
- 765 Lagerquist, R., McGovern, A., & Gagne, D. J., II. (2019, August). Deep learning for  
766 spatially explicit prediction of Synoptic-Scale fronts. *Weather Forecast.*, *34*(4),  
767 1137–1160. doi: 10.1175/WAF-D-18-0183.1
- 768 Lapuschkin, S. (2019). *Opening the machine learning black box with layer-wise*  
769 *relevance propagation* (Doctoral dissertation, Technischen Universität Berlin,  
770 Berlin, Germany). doi: 10.14279/DEPOSITONCE-7942
- 771 Lee, Y., Kummerow, C. D., & Ebert-Uphoff, I. (2021, April). Applying machine  
772 learning methods to detect convection using geostationary operational environ-  
773 mental satellite-16 (GOES-16) advanced baseline imager (ABI) data. *Atmos.*  
774 *Meas. Tech.*, *14*(4), 2699–2716. doi: 10.5194/amt-14-2699-2021
- 775 Li, J., Thompson, D. W. J., Barnes, E. A., & Solomon, S. (2017, December). Quan-  
776 tifying the lead time required for a linear trend to emerge from natural climate  
777 variability. *J. Clim.*, *30*(24), 10179–10191. doi: 10.1175/JCLI-D-16-0280.1
- 778 Lu, E., Chen, H., Tu, J., Song, J., Zou, X., Zhou, B., ... Jiang, Z. (2015, De-  
779 cember). The nonlinear relationship between summer precipitation in  
780 china and the sea surface temperature in preceding seasons: A statistical  
781 demonstration. *J. Geophys. Res. D: Atmos.*, *120*(23), 12,027–12,036. doi:  
782 10.1002/2015JD024030
- 783 Madakumbura, G. D., Thackeray, C. W., Norris, J., Goldenson, N., & Hall, A.

- 784 (2021, July). Anthropogenic influence on extreme precipitation over global  
 785 land areas seen in multiple observational datasets. *Nat. Commun.*, *12*(1), 3944.  
 786 doi: 10.1038/s41467-021-24262-x
- 787 Maher, N., Milinski, S., & Ludwig, R. (2021). Large ensemble climate model simula-  
 788 tions: introduction, overview, and future prospects for utilising multiple types  
 789 of large ensemble. *Earth System Dynamics*, *12*(2), 401–418.
- 790 Mahony, C. R., & Cannon, A. J. (2018, February). Wetter summers can intensify  
 791 departures from natural variability in a warming climate. *Nat. Commun.*, *9*(1),  
 792 783. doi: 10.1038/s41467-018-03132-z
- 793 Mamalakis, A., Ebert-Uphoff, I., & Barnes, E. A. (2021). Neural network attribu-  
 794 tion methods for problems in geoscience: A novel synthetic benchmark dataset.  
 795 *arXiv*. Retrieved from <http://arxiv.org/abs/2103.10005>
- 796 Mankin, J. S., Lehner, F., Coats, S., & McKinnon, K. A. (2020, October). The  
 797 value of initial condition large ensembles to robust adaptation decision-making.  
 798 *Earths Future*, *8*(10). doi: 10.1029/2020ef001610
- 799 Marvel, K., & Bonfils, C. (2013, November). Identifying external influences on global  
 800 precipitation. *Proc. Natl. Acad. Sci. U. S. A.*, *110*(48), 19301–19306. doi: 10  
 801 .1073/pnas.1314382110
- 802 Montavon, G., Binder, A., Lapuschkin, S., Samek, W., & Müller, K.-R. (2019).  
 803 Layer-Wise relevance propagation: An overview. In W. Samek, G. Montavon,  
 804 A. Vedaldi, L. K. Hansen, & K.-R. Müller (Eds.), *Explainable AI: Interpreting,*  
 805 *explaining and visualizing deep learning* (Vol. 11700, pp. 193–209). Cham:  
 806 Springer International Publishing. doi: 10.1007/978-3-030-28954-6\\_10
- 807 Mudelsee, M. (2019, March). Trend analysis of climate time series: A review of  
 808 methods. *Earth-Sci. Rev.*, *190*, 310–322. doi: 10.1016/j.earscirev.2018.12.005
- 809 Newell, G. J., & Lee, B. (1981, May). Ridge regression: An alternative to multiple  
 810 linear regression for highly correlated data. *J. Food Sci.*, *46*(3), 968–969. doi:  
 811 10.1111/j.1365-2621.1981.tb15400.x
- 812 North, G. R., & Stevens, M. J. (1998, April). Detecting climate signals in  
 813 the surface temperature record. *J. Clim.*, *11*(4), 563–577. doi: 10.1175/  
 814 1520-0442(1998)011<0563:DCSITS>2.0.CO;2
- 815 O’Gorman, P. A. (2015). Precipitation extremes under climate change. *Curr Clim*  
 816 *Change Rep*, *1*(2), 49–59. doi: 10.1007/s40641-015-0009-3
- 817 O’Neill, B. C., Tebaldi, C., van Vuuren, D. P., Eyring, V., Friedlingstein, P., Hurtt,  
 818 G., . . . Sanderson, B. M. (2016, September). The scenario model intercompari-  
 819 son project (ScenarioMIP) for CMIP6. *Geoscientific Model Development*, *9*(9),  
 820 3461–3482. doi: 10.5194/gmd-9-3461-2016

- 821 Rohde, R. A., & Hausfather, Z. (2020, December). The Berkeley Earth land/ocean  
 822 temperature record. *Earth Syst. Sci. Data*, *12*(4), 3469–3479. doi: 10.5194/  
 823 essd-12-3469-2020
- 824 Rosenzweig, C., Tubiello, F. N., Goldberg, R., Mills, E., & Bloomfield, J. (2002, Oc-  
 825 tober). Increased crop damage in the US from excess precipitation under cli-  
 826 mate change. *Glob. Environ. Change*, *12*(3), 197–202. doi: 10.1016/S0959  
 827 -3780(02)00008-0
- 828 Rudin, C. (2019, May). Stop explaining black box machine learning models for high  
 829 stakes decisions and use interpretable models instead. *Nature Machine Intelli-*  
 830 *gence*, *1*(5), 206–215. doi: 10.1038/s42256-019-0048-x
- 831 Sanderson, B. M., Oleson, K. W., Strand, W. G., Lehner, F., & O’Neill, B. C. (2018,  
 832 February). A new ensemble of GCM simulations to assess avoided impacts in  
 833 a climate mitigation scenario. *Clim. Change*, *146*(3), 303–318. doi: 10.1007/  
 834 s10584-015-1567-z
- 835 Santer, B. D., Fyfe, J. C., Solomon, S., Painter, J. F., Bonfils, C., Pallotta, G., &  
 836 Zelinka, M. D. (2019, October). Quantifying stochastic uncertainty in detec-  
 837 tion time of human-caused climate signals. *Proc. Natl. Acad. Sci. U. S. A.*,  
 838 *116*(40), 19821–19827. doi: 10.1073/pnas.1904586116
- 839 Santer, B. D., Mears, C., Doutriaux, C., Caldwell, P., Gleckler, P. J., Wigley,  
 840 T. M. L., . . . Wentz, F. J. (2011, November). Separating signal and noise  
 841 in atmospheric temperature changes: The importance of timescale. *J. Geophys.*  
 842 *Res.*, *116*(D22). doi: 10.1029/2011jd016263
- 843 Santer, B. D., Taylor, K. E., Wigley, T. M. L., Johns, T. C., Jones, P. D., Karoly,  
 844 D. J., . . . Tett, S. (1996, July). A search for human influences on the  
 845 thermal structure of the atmosphere. *Nature*, *382*(6586), 39–46. doi:  
 846 10.1038/382039a0
- 847 Scaife, A. A., & Smith, D. (2018, July). A signal-to-noise paradox in climate science.  
 848 *npj Climate and Atmospheric Science*, *1*(1), 1–8. doi: 10.1038/s41612-018-0038  
 849 -4
- 850 Schneider, T., & Held, I. M. (2001, February). Discriminants of Twentieth-Century  
 851 changes in earth surface temperatures. *J. Clim.*, *14*(3), 249–254. doi: 10.1175/  
 852 1520-0442(2001)014(0249:LDOTCC)2.0.CO;2
- 853 Schulzweida, U. (2019). CDO user guide (version 1.9. 6). *Max Planck Institute for*  
 854 *Meteorology: Hamburg, Germany.*
- 855 Silva, S. J., Ma, P.-L., Hardin, J. C., & Rothenberg, D. (2021, May). Physically reg-  
 856 ularized machine learning emulators of aerosol activation. *Geosci. Model Dev.*,  
 857 *14*(5), 3067–3077. doi: 10.5194/gmd-14-3067-2021

- 858 Sippel, S., Meinshausen, N., Fischer, E. M., Székely, E., & Knutti, R. (2020, Jan-  
 859 uary). Climate change now detectable from any single day of weather at global  
 860 scale. *Nat. Clim. Chang.*, *10*(1), 35–41. doi: 10.1038/s41558-019-0666-7
- 861 Solomon, A., & Newman, M. (2012, July). Reconciling disparate twentieth-century  
 862 Indo-Pacific ocean temperature trends in the instrumental record. *Nat. Clim.*  
 863 *Chang.*, *2*(9), 691–699. doi: 10.1038/nclimate1591
- 864 Solow, A. R. (1987, October). Testing for climate change: An application of the  
 865 Two-Phase regression model. *J. Appl. Meteorol. Climatol.*, *26*(10), 1401–1405.  
 866 doi: 10.1175/1520-0450(1987)026<1401:TFCCAA>2.0.CO;2
- 867 Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R.  
 868 (2014). Dropout: a simple way to prevent neural networks from overfitting.  
 869 *J. Mach. Learn. Res.*, *15*(1), 1929–1958.
- 870 Swain, D. L., Wing, O. E. J., Bates, P. D., Done, J. M., Johnson, K. A., &  
 871 Cameron, D. R. (2020, November). Increased flood exposure due to climate  
 872 change and population growth in the united states. *Earths Future*, *8*(11). doi:  
 873 10.1029/2020ef001778
- 874 Tabari, H., & Willems, P. (2018, March). Seasonally varying footprint of climate  
 875 change on precipitation in the middle east. *Sci. Rep.*, *8*(1), 4435. doi: 10.1038/  
 876 s41598-018-22795-8
- 877 Toms, B. A., Barnes, E. A., & Ebert-Uphoff, I. (2020, September). Physically  
 878 interpretable neural networks for the geosciences: Applications to earth sys-  
 879 tem variability. *J. Adv. Model. Earth Syst.*, *12*(9), e2019MS002002. doi:  
 880 10.1029/2019ms002002
- 881 USGCRP. (2018). *Impacts, risks, and adaptation in the united states: Fourth na-*  
 882 *tional climate assessment, volume II* (Tech. Rep.). Washington, DC, USA:  
 883 U.S. Global Change Research Program. doi: 10.7930/NCA4.2018
- 884 Weyn, J. A., Durran, D. R., & Caruana, R. (2020, September). Improving data-  
 885 driven global weather prediction using deep convolutional neural networks  
 886 on a cubed sphere. *J. Adv. Model. Earth Syst.*, *12*(9), e2020MS002109. doi:  
 887 10.1029/2020ms002109
- 888 Wills, R. C., Battisti, D. S., Armour, K. C., Schneider, T., & Deser, C. (2020, Oc-  
 889 tober). Pattern recognition methods to separate forced responses from inter-  
 890 nal variability in climate model ensembles and observations. *J. Clim.*, *33*(20),  
 891 8693–8719. doi: 10.1175/JCLI-D-19-0855.1
- 892 Wills, R. C., Schneider, T., Wallace, J. M., Battisti, D. S., & Hartmann, D. L.  
 893 (2018, March). Disentangling global warming, multidecadal variability, and  
 894 el niño in pacific temperatures. *Geophys. Res. Lett.*, *45*(5), 2487–2496. doi:

895 10.1002/2017GL076327

896 Wu, J., Chen, X., Yao, H., Gao, L., Chen, Y., & Liu, M. (2017, August). Non-  
897 linear relationship of hydrological drought responding to meteorological  
898 drought and impact of a large reservoir. *J. Hydrol.*, *551*, 495–507. doi:  
899 10.1016/j.jhydrol.2017.06.029

900 Zadeh, L. A. (1965). Information and control. *Fuzzy Sets and Systems*, *8*(3), 338–  
901 353. doi: 10.1002/joc.1027

902 Zappa, G., Hoskins, B. J., & Shepherd, T. G. (2015, August). Improving climate  
903 change detection through optimal seasonal averaging: The case of the north  
904 atlantic jet and european precipitation. *J. Clim.*, *28*(16), 6381–6397. doi:  
905 10.1175/JCLI-D-14-00823.1