

1 ***This Looks Like That There: Interpretable neural networks for image tasks***
2 **when location matters**

3 Elizabeth A. Barnes,^a Randal J. Barnes,^b Zane K. Martin,^a and Jamin K. Rader,^a

4 ^a*Department of Atmospheric Science, Colorado State University, Fort Collins, CO, USA.*

5 ^b*Civil, Environmental, and Geo- Engineering, University of Minnesota, Minneapolis, MN, USA.*

6 *Corresponding author:* Elizabeth A. Barnes, eabarnes@colostate.edu

7 ABSTRACT: We develop and demonstrate a new interpretable deep learning model specifically
8 designed for image analysis in earth system science applications. The neural network is designed to
9 be inherently interpretable, rather than explained via *post hoc* methods. This is achieved by training
10 the network to identify parts of training images that act as prototypes for correctly classifying
11 unseen images. The new network architecture extends the interpretable prototype architecture of a
12 previous study in computer science to incorporate absolute location. This is useful for earth system
13 science where images are typically the result of physics-based processes, and the information is
14 often geo-located. Although the network is constrained to only learn via similarities to a small
15 number of learned prototypes, it can be trained to exhibit only a minimal reduction in accuracy
16 compared to non-interpretable architectures. We apply the new model to two earth science use
17 cases: a synthetic data set that loosely represents atmospheric high- and low-pressure systems, and
18 atmospheric reanalysis fields to identify the state of tropical convective activity associated with
19 the Madden-Julian oscillation. In both cases, we demonstrate that considering absolute location
20 greatly improves testing accuracies when compared to a location-agnostic method. Furthermore,
21 the network architecture identifies specific historical dates that capture multivariate, prototypical
22 behaviour of tropical climate variability.

23 SIGNIFICANCE STATEMENT: Machine learning models are incredibly powerful predictors,
24 but are often opaque “black boxes”. The how-and-why the model makes its predictions is inscrutable
25 — the model is not interpretable. We introduce a new machine learning model specifically designed
26 for image analysis in earth system science applications. The model is designed to be inherently
27 interpretable and extends previous work in computer science to incorporate location information.
28 This is important because images in earth system science are typically the result of physics-based
29 processes, and the information is often map based. We demonstrate its use for two earth science use
30 cases and show that the interpretable network exhibits only a small reduction in accuracy compared
31 to black box models.

32 1. Introduction

33 Machine learning has been identified as an innovative, under-explored tool for furthering under-
34 standing and simulation of the Earth system (Balmaseda et al. 2020; Irrgang et al. 2021; National
35 Academies of Sciences Engineering and Medicine 2020). Artificial neural networks (as a type of
36 supervised machine learning) have emerged as a powerful tool for extracting nonlinear relationships
37 amidst noisy data, and thus are particularly suited to this endeavor. However, a major criticism of
38 the use of neural network models for scientific applications is that they are “black boxes.” Scientists
39 typically want to know why the model reached the decision that it did. The benefit of explaining
40 the decision-making process of a model goes beyond that of satisfying curiosity: explanation can
41 assist users in (1) determining if the model is getting the right answers for the right reasons (e.g.
42 Lapuschkin et al. 2019), (2) controlling and improving the machine learning approach (e.g. Keys
43 et al. 2021), and (3) discovering new science (e.g. Toms et al. 2020; Barnes et al. 2020). Effective
44 explanations also increase user confidence.

45 Because researchers are driven by the desire to explain the decision-making process of deep
46 learning models, a large variety of *post hoc* explainability methods have been developed (e.g.
47 Buhrmester et al. 2019; Barredo Arrieta et al. 2020; Samek et al. 2021). By *post hoc*, we mean
48 methods in which a deep learning model has already been trained and the user attempts to explain
49 the predictions of the black box model after the predictions have been made. Although *post hoc*
50 explainability methods have demonstrated success across many scientific applications (including
51 earth system science, e.g. McGovern et al. (2019); Toms et al. (2020); Davenport and Diffenbaugh

52 (2021)), they are not without their drawbacks. *Post hoc* explainability methods do not exactly
53 replicate the computations made by the black box model. Instead, through a set of assumptions and
54 simplifications, these methods quantify some reduced version of the model (e.g. Montavon et al.
55 2018) and, thus, do not explain the actual decision-making process of the network. Furthermore,
56 the explanations are not always reliable (Kindermans et al. 2019). Different explanation methods
57 can produce vastly different explanations of the exact same black box model (Mamalakis et al. 2021,
58 2022). Even if the explanation is reliable, at times the output of the explainability method itself
59 requires extensive deciphering by the scientist to understand the result (e.g. Mayer and Barnes
60 2021; Martin et al. 2021; Barnes et al. 2020). Rudin (2019) discusses in detail many of these
61 potential issues with explainable machine learning methods and suggests that we should instead
62 be using machine learning models that are inherently “interpretable”. That is, instead of trying
63 to explain black box models, we should be creating models where the decision-making process is
64 interpretable by design.

65 Chen et al. (2019) present an example of one type of interpretable neural network, the *prototypical*
66 *part network* (ProtoPNet). The ProtoPNet hinges on training a neural network to identify patches
67 of the training images that act as “prototypes” for correctly classifying unseen images. The idea
68 for the ProtoPNet stems from the need to define a form of interpretability that works the way a
69 scientist might describe their way of thinking. In their specific application, Chen et al. (2019)
70 focus on classifying images of birds by their species. A scientist may classify a new bird image
71 by comparing it to representative examples of each species (i.e. species prototypes) and choosing
72 the prototype that most resembles the image, i.e. *this looks like that*. In this way, the network is
73 inherently interpretable in that the *actual* decision-making process can be linked to specific features
74 of the bird in the input image and their similarity to a relatively small number of species-specific
75 prototypes that are directly drawn from the training set. For bird species identification, Chen et al.
76 (2019) demonstrate that the ProtoPNet learns prototypes that represent distinguishing features such
77 as the red head of a red-bellied woodpecker, or the bright blue wing of a Florida jay.

78 Images in earth system science are typically the result of physics-based processes, and the
79 information is often geo-located. Thus, unlike the ProtoPNet of Chen et al. (2019) which does not
80 care *where* the bird’s wing is in the image, the location of specific earth system features can be
81 critical to the final task (although this is certainly not always the case, e.g. identification of cloud

82 types from satellite imagery; Rasp et al. (2019)). For example, the mere presence of a low-pressure
83 system on a weather map is not enough to know where it will rain. Instead, the location of the low
84 — *where* it is — is also vital for this task. Similarly, identifying the presence of a strong El Niño
85 requires not only warm sea-surface temperatures, but specifically warm sea-surface temperatures
86 in the tropical equatorial east Pacific (e.g. Philander 1983). Here, we extend the ProtoPNet of Chen
87 et al. (2019) to consider absolute location in the interpretable prototype architecture, which we
88 call the ProtoLNet (“Prototypical Location Network”). We demonstrate that considering absolute
89 location greatly improves the network accuracy (ProtoLNet rather than ProtoPNet) for two earth
90 science use cases. The first use case, the idealized quadrants use case (Section 3), applies the
91 ProtoLNet to a synthetic data set that loosely represents high- and low-pressure systems where the
92 need for location information is readily apparent. The second use case applies the ProtoLNet to
93 over 100 years of atmospheric reanalysis fields to identify the state of tropical convective activity
94 associated with the Madden-Julian oscillation (MJO; Madden and Julian 1971, 1972; Zhang 2005).
95 The MJO use case (Section 4) provides a real, geophysical example of how the ProtoLNet relies
96 on location information to make its predictions and demonstrates how the learned prototypes can
97 be viewed as prototypical behaviour of transient climate phenomena.

98 **2. Network Design & Training**

99 As discussed in the introduction, the ProtoLNet is largely based on the ProtoPNet of Chen et al.
100 (2019). We describe the network architecture below, highlighting where our ProtoLNet diverges
101 from the ProtoPNet of Chen et al. (2019). We then describe the training procedure in detail.

102 *a. ProtoLNet architecture*

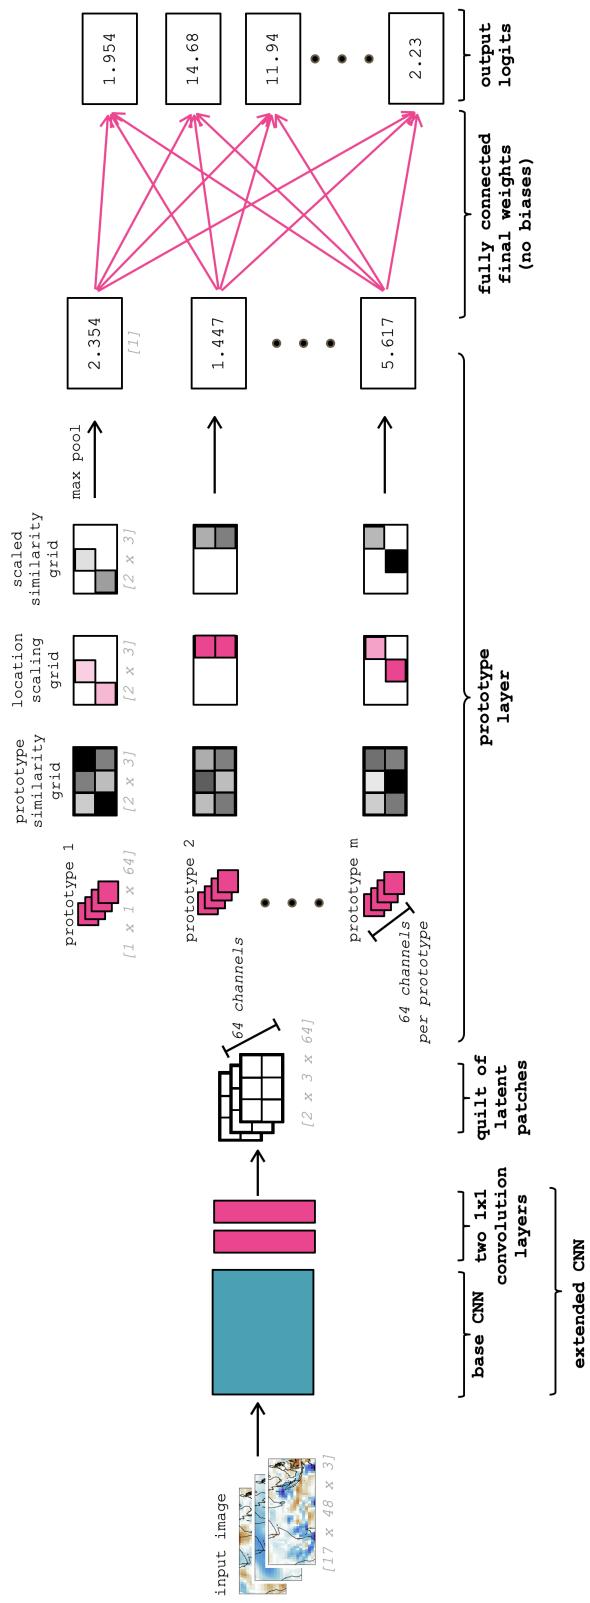
103 The ProtoLNet is designed to classify images by comparing latent patches of the input image
104 to prototypical latent patches learned from the training set, all while explicitly considering the
105 location within the image of the similar latent patches. Throughout, we use the word “patch”
106 to refer to a group of neighboring pixels within the input image, and “latent patch” to refer to a
107 latent representation of a patch that is computed via a series of convolutional and pooling layers
108 within the convolutional neural network. In this section, we first provide a general overview of

109 the ProtoLNet architecture from start to finish, and then go into more detail about each step in
110 subsequent paragraphs, ending with the training process.

114 The ProtoLNet architecture (Fig. 1) is very similar to that of the ProtoPNet, and starts with a
115 base convolutional neural network (CNN) chosen by the user that takes-in an image as input. As
116 discussed more in Section c, this base CNN may be a pre-trained network, or a newly initialized
117 network with randomized weights. The CNN is followed by two 1×1 convolutional layers that act
118 to restructure the dimensions of the CNN output to be consistent with the subsequent prototype
119 layer. It is within the prototype layer that the interpretable learning is done. The network is
120 trained to learn representative latent patches within the training set specific to each class, termed
121 *prototypes*, which provide evidence for the image belonging to a particular class. That is, when
122 the input image has a patch whose latent representation *looks like that* prototype, it is labeled as
123 belonging to the prototype’s associated class. This is done by computing the similarity of each
124 prototype to the latent patches of the input image. Unique to our ProtoLNet, these similarity scores
125 are scaled by a learned, prototype-specific location scaling grid so that similarities to the prototypes
126 are only important for certain locations within the input image. The maximum scaled similarity
127 score across the latent patches for each prototype is then computed. These scores are connected to
128 the output via a fully connected layer, and the weighted scores are summed for each output class to
129 produce a total number of “points” for each class. The class with the highest number of points is
130 then identified as the predicted class.

131 As will be discussed in detail in Section c, the ProtoLNet learns the convolutional kernels within
132 the two 1×1 convolution layers, the prototypes, the location scaling grid, and the final fully
133 connected weights (pink components in Fig. 1). The user must specify the number of prototypes
134 specific to each output class. For the use cases presented here, we choose an equal number of
135 prototypes for each class, so if there are n classes and p prototypes per class, then there are $m = n * p$
136 total prototypes. A critical aspect of the architecture is that *each prototype is assigned to only one*
137 *class* since it is used as evidence that a particular sample belongs its class.

138 Each sample is pushed through the extended CNN, which results in an output “quilt” of latent
139 patches. To introduce some general notation, the quilt has shape $a \times b \times D$, where $a \times b$ is the
140 new image shape after undergoing pooling in the base CNN, and D corresponds to the number of
141 convolutional kernels chosen by the user. Each prototype vector (\mathbf{p}) then has shape $1 \times 1 \times D$. To



111 Fig. 1. Schematic depicting the ProtoNet architecture. Example and internally consistent dimensions of the tensors at each step are given in grey
 112 brackets, although the specific dimensions vary for each use case. Pink colors denote components of the network that are trained (learned), while grey
 113 and black colors denote components that are directly computed. The weights within the base CNN (blue shading) can either be trained or frozen.

142 simplify our discussion, from here forward we will drop the general notation and instead use the
 143 specific dimensions (denoted in gray) of the example shown in Fig. 1. That is, $a = 2$, $b = 3$, and
 144 $D = 64$.

145 For the example in Fig. 1, a latent patch has shape $1 \times 1 \times 64$, and the quilt of latent patches
 146 output by the extended CNN has shape $2 \times 3 \times 64$. Because the input image has already potentially
 147 undergone multiple convolutional and pooling layers within the extended CNN, these latent patches
 148 *do not* represent a single pixel of the input image, but instead are a latent representation of some
 149 larger patch within the input image. Similar to the latent patches, each of the m learned prototypes
 150 are a latent representation of some larger region of the input image. Each prototype has the same
 151 shape as a latent patch: $1 \times 1 \times 64$. The similarity score for a prototype \mathbf{p} and a latent patch \mathbf{z}
 152 is computed as a function of the distance between these two vectors (i.e. the L_2 norm of the
 153 difference). The greater the distance between, the lower the similarity score. Following Chen et al.
 154 (2019), we compute

$$155 \quad \text{SimilarityScore} = \log \left(\frac{\|\mathbf{z} - \mathbf{p}\|_2^2 + 1}{\|\mathbf{z} - \mathbf{p}\|_2^2 + \epsilon} \right) \approx \log \left(1 + \frac{1}{(\text{distance})^2} \right) \quad (1)$$

156 where $\|\cdot\|_2^2$ is the squared L_2 norm and ϵ is a small number, there to guard against divide-by-zero
 157 problems. Applying this similarity metric to a quilt of latent patches results in $m \times 2 \times 3$ similarity
 158 grids, one for each prototype. The values within these grids thus quantify how much that latent
 159 patch of the input *looks like* each prototype.

160 In the original ProtoPNet, at this point the maximum similarity within each similarity grid is
 161 computed for each prototype. However, unique to our ProtoLNet — and indeed the novelty of
 162 this work — is that we scale each prototype’s similarity grid by a location-specific value learned
 163 by the network. This step rescales the similarities such that similarities in certain locations are
 164 accentuated and similarities in other locations are muted. To follow this paper’s title, it isn’t enough
 165 for *this* latent patch (at any location) to look like *that* prototype. Instead, *this* latent patch must look
 166 like *that* prototype in only specific locations — *there*. This results in m location-scaled similarity
 167 grids, one for each prototype.

168 Once again following the architecture of the original ProtoPNet, we apply max pooling to each
 169 scaled similarity grid to obtain a single score for the maximum similarity (scaled by the location

170 scaling) between a prototype and the input image. These scores are then connected to the output
171 layer via a fully connected layer with learned weights but zero bias. The choice of zero bias in the
172 final fully-connected layer is essential for interpreting the prototypes as providing evidence for a
173 particular class. With a zero bias, the final points contributing to each class are comprised only of
174 a sum of location-scaled similarity scores multiplied by a final weight. The final weights layer is
175 trained separately from the rest of the network. The layer is trained in such a way as to keep weights
176 connecting prototypes with their associated class large, while minimizing the weights connecting
177 prototypes with their non-class output units (see Section c). Finally, as is standard with a fully
178 connected layer, the output values (weighted scores) contributing to each output unit are summed
179 to produce a total number of points for each class. The class with the highest number of points is
180 identified as the predicted class.

181 In the original ProtoPNet, there was no location scaling. Without this location scaling, the
182 network is agnostic to *where* the input image looks most like each prototype. That is, the only
183 thing of import is that the image looks like the prototype *somewhere*. Returning to the example of
184 classifying bird images (as explored in Chen et al. (2019)), a prototype may correspond to a latent
185 representation of the red head of a red-bellied woodpecker. The original ProtoPNet does not care
186 whether a red head is found in the upper left or the upper right of the input image. Rather, the
187 ProtoPNet just considers whether a red head is present at all. For our ProtoLNet presented here,
188 the network is designed to take into consideration not only that a red head is found, but also *where*
189 within the image the red head occurs. As we will show, this consideration of location can be highly
190 beneficial in geophysical applications.

191 *b. Choosing the base CNN*

192 We envision three main approaches to choosing a base CNN. The first takes an existing CNN that
193 has been previously trained to perform classification tasks. This CNN may already be performing
194 well, but interpretability is desired. The user removes the output layer and fully connected layers of
195 their existing CNN and then use the result as their base CNN for the ProtoLNet. In this approach,
196 the ProtoLNet is used purely for interpretability of the original CNN.

197 The second approach to choosing a base CNN is to, once again, take a pre-trained CNN, remove
198 the output and fully connected layers, and then use the result as the base CNN for the ProtoLNet.

Stage	Type	Base CNN	1x1 Layers	Prototypes	Location Scaling	Final Weights
1	train prototypes	frozen/train	train	train	train	frozen
2	replace prototypes	frozen	frozen	replace	frozen	frozen
3	train weights	frozen	frozen	frozen	frozen	train

FIG. 2. The three different stages of training the ProtoLNet.

199 The difference is that now the user allows the weights within the base CNN to be further refined
200 during the ProtoLNet training in order to optimize the performance of the ProtoLNet. Allowing
201 the base CNN weights to be updated implies that the user is no longer interpreting the same base
202 CNN with which they started. However, if the goal is to create an interpretable network that is
203 as accurate as possible, this may be a good approach. Furthermore, for image classification tasks,
204 one might choose to use a CNN previously trained on a large dataset, e.g. VGG-19 (Simonyan and
205 Zisserman 2014), as done by Chen et al. (2019).

206 The third approach to choosing a base CNN applies when no suitable pre-trained base CNN
207 exists. In this case, the user must train the interpretable network from scratch. In this instance,
208 there are two main choices. A separate base CNN could be trained, stripped of its final output and
209 fully connected layers, and then appended to the ProtoLNet (as discussed above). Alternatively,
210 one could initialize the base CNN with random initial weights and train it directly within the
211 ProtoLNet architecture. We have tried both methods for the use cases explored here and found
212 that they produced similar accuracies (although we acknowledge this may not always be the case).
213 Here, we present results where we first pre-train a base CNN and then append it to the ProtoLNet,
214 in order to provide a base accuracy with which to compare our ProtoLNet results.

215 *c. ProtoLNet training*

216 The training of the ProtoLNet is done in triads of stages (Fig. 2), largely following the original
217 training approach of Chen et al. (2019). The first stage of training involves learning the prototypes
218 by training the 1×1 layers, prototypes, location scaling grid, and the base CNN (if desired by the
219 user; see Section b) at the same time. The final weights are frozen during this stage. The second
220 stage of training involves replacing each prototype with the nearest latent patch within the training

221 samples of the same class. That is, stage 1 allows the network to learn *any* form of the prototype
 222 latent patch, and stage 2 replaces this prototype with the most similar training latent patch from
 223 the same class. In this way, the prototypes always directly correspond to a latent patch in one
 224 particular training sample. In the third stage of training, we freeze all elements of the ProtoLNet
 225 except for the fully connected final weights (pink arrows in Fig. 1), and the network learns them
 226 alone. These three stages are then cycled through multiple times (for our use cases, up to 5 times)
 227 for full training of the ProtoLNet.

228 *Initialization:* Prior to stage 1, the two 1×1 convolutional layers are initialized with random
 229 values drawn from a truncated normal distribution (He et al. 2015). The prototypes are initialized
 230 with random values drawn from a uniform distribution between 0.0 and 1.0, and the location scaling
 231 grid is initialized with ones everywhere (see Appendix B for additional details). The final weights
 232 (w) that connect a prototype with its assigned class are given an initial value of 1.0, and all other
 233 final weights are initialized to -0.5. The initialization of the base CNN was already discussed in
 234 Section b.

235 *Stage 1:* Training is performed via stochastic gradient descent with the Adam optimizer and
 236 batch size of 32. For the quadrants use case, the learning rate is set to 0.01 for every stage 1
 237 cycle. For the MJO use case, the learning rate is also initially set to 0.01 but is reduced by an
 238 order of magnitude for the third cycle of stage 1 and every cycle thereafter. The network is trained
 239 with the standard cross-entropy loss (e.g. Géron 2019) added to two additional loss terms: the
 240 ClusterCost and SeparationCost. The cross-entropy loss penalizes the network for misclassifying
 241 the training samples. The ClusterCost encourages the network to construct prototypes such that
 242 training images have at least one latent patch with high similarity to a prototype of the correct
 243 class. The SeparationCost discourages the network from constructing prototypes such that training
 244 images have any latent patches with a high similarity to prototypes of the incorrect classes. Thus,
 245 the full stage 1 loss function takes the form

$$246 \quad \text{Loss} = \text{CrossEntropy} + \beta_1 \text{ClusterCost} - \beta_2 \text{SeparationCost} \quad (2)$$

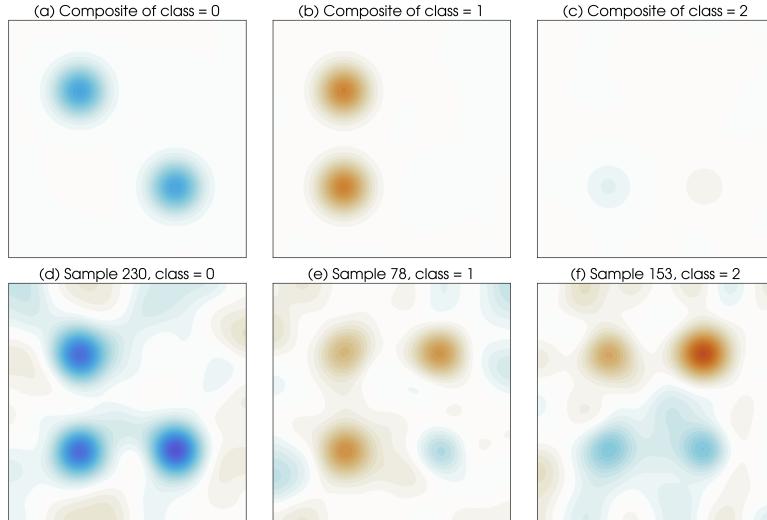
247 where β_1 and β_2 are coefficients chosen by the user. Full forms of the ClusterCost and SeparationCost,
 248 along with their coefficient values, are provided in Appendix C. For all use cases, we
 249 train in stage 1 for 10 epochs before moving to stage 2 of training.

250 *Stage 2:* This stage does not involve any iterative training but instead is direct computation.
251 Specifically, the similarity scores are computed between each learned prototype from stage 1 and
252 every latent patch of every training image of the same class. The prototype is then *replaced* by
253 the training latent patch with the highest similarity. Note that this replacement process will nearly
254 always reduce the accuracy of the network because it replaces the stage 1-optimized prototypes
255 with something from the training set. However, this step is central to the interpretability of the
256 ProtoLNet. By cycling through all three training stages multiple times, the network learns to
257 perform well using the replaced prototypes from the training set.

258 *Stage 3:* The final weights $w_{k,j}$ connecting prototypes of class k to the output class j are learned
259 via convex optimization, since all other layers are frozen. As a reminder, all $w_{k,j}$ for $k = j$ are
260 initialized to 1.0, and the rest, $w_{k,j}$ for $k \neq j$, are initialized to -0.5. The weights are frozen for
261 stages 1 and 2 of training. In stage 3, all other free parameters in the ProtoLNet are frozen, and the
262 weights alone are trained to minimize the cross-entropy loss of the final output plus an additional
263 L_1 regularization term evaluated on the weights $w_{k,j}$ for $k \neq j$. This additional loss term provides
264 sparsity to the final model, i.e. $w_{k,j} \approx 0$ for $k \neq j$, which reduces the use of negative reasoning by
265 the network (“this does *not* look like that”). See Singh and Yow (2021) for an exploration of the
266 consequences when this sparsity requirement is relaxed. For the idealized quadrants use case, we
267 set the regularization parameter to 0.5. For the MJO use case, it is set to 0.1. For all use cases, we
268 train in stage 3 for 10 epochs. At that point, we either end training completely (i.e. we have the
269 fully trained ProtoLNet), or we cycle through stages 1-3 again.

270 **3. Use Case: Idealized Quadrants**

271 As a first demonstration of the ProtoLNet, we construct an idealized synthetic test set to loosely
272 represent the horizontal (latitude by longitude) spatial structures of geophysical anomalies. For
273 example, the synthetic fields (or images) could represent idealized low- and high-pressure circu-
274 lations. The anomaly fields are 100x100 pixels in size and are constructed by first initializing the
275 field with random Gaussian noise. We then randomly add an additional anomaly value (uniformly
276 distributed between 2 and 15) to the center of one or more of the four quadrants of each square
277 field. Finally, we smooth each field with a Gaussian filter with standard deviation of 7 to make



280 FIG. 3. The top three panels (a-c) show composites of all samples by class label for the idealized quadrants
 281 use case. The bottom three panels (d-f) exhibit one example sample for each class.

278 the fields look more like typical tropospheric pressure anomalies. Example samples are shown in
 279 Fig. 3.

282 The fields in the idealized data set are assigned labels based on the sign of the anomalies in each
 283 of the four quadrants of the sample (Fig. 3). Specifically, fields with negative anomalies in both the
 284 second and fourth quadrants are labeled class 0, fields with positive anomalies in both the second
 285 and third quadrants are labeled class 1, and all other fields are labeled class 2 (Fig. 3a-c). Fig. 3d-f
 286 show example samples for each class. As designed, sample #230 (labeled class 0) has negative
 287 anomalies in the second and fourth quadrants, sample #78 (labeled class 1) has positive anomalies
 288 in the second and third quadrants, and sample #153 (labeled class 2) does not achieve either of the
 289 requirements of classes 0 or 1. As will become clear, this idealized data set was designed such that
 290 the location of the different anomalies matters.

291 The synthetic data set has equally balanced classes by construction, with 3,000 samples for each
 292 of the three classes (9,000 samples total). This set is then randomly split such that 7,200 samples
 293 are used for training and 1,800 for testing. Prior to training, the input images are standardized by
 294 subtracting the mean and dividing by the standard deviation over all training pixels.

295 We task the ProtoLNet with ingesting a single input field and classifying it into one of the
 296 three classes, as depicted in Fig. 4. The network cannot simply identify the existence of negative

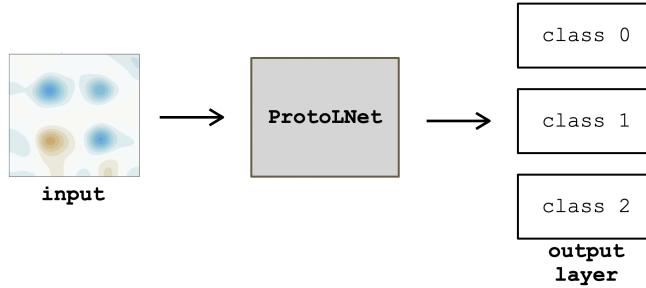
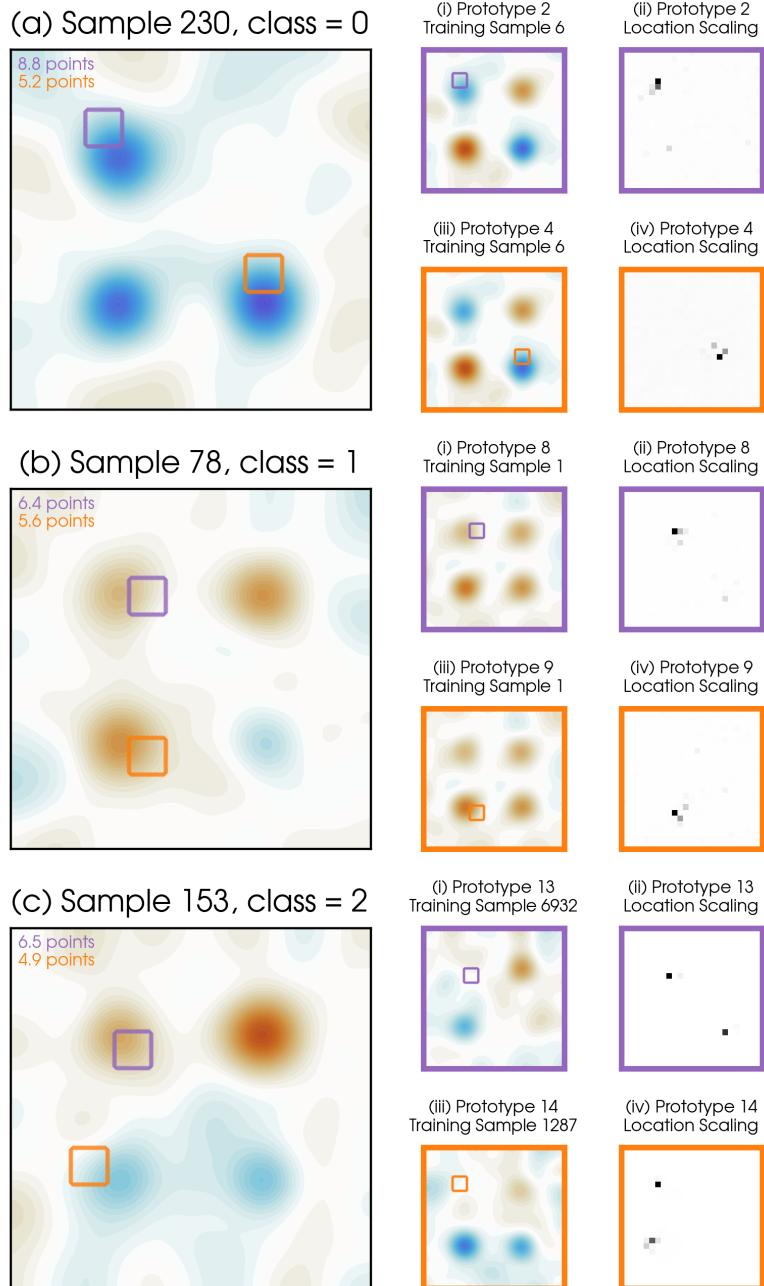


FIG. 4. Prediction setup for the idealized quadrants use case.

297 anomalies (in the case of class 0) or the existence of positive anomalies (in the case of class 1).
 298 Instead, it must consider the existence of different signed anomalies *and their location* within the
 299 input field. To illustrate this point, we trained a ProtoPNet where location is not considered (i.e.
 300 learning of the location scaling grid is turned off) and — unsurprisingly — the network fails with
 301 an accuracy of 32%, no better than random chance (i.e. 33%).

302 We first train a standard CNN to perform the classification task and act as our base CNN for
 303 the ProtoLNet. Details of the CNN architecture and training parameters are provided in Appendix
 304 A. Once the CNN is trained, we remove the final fully connected layer and output layer, and
 305 append the result to the ProtoLNet to become the base CNN (see Fig. 1). We assign 5 prototypes
 306 (with $D = 128$) to each output class, for a total of 15 prototypes. Using more prototypes than
 307 this yielded prototypes that rarely provided points for any sample. We cycle through the three
 308 stages of ProtoLNet training (Fig. 2) five times, freezing the base CNN for the first cycle of stage
 309 1 but allowing it to train for all subsequent cycles of stage 1. Once fully trained, the ProtoLNet
 310 achieves an accuracy of 96%, a significant improvement over random chance and the ProtoPNet.
 311 For comparison, the base CNN achieves an accuracy of 98%. The ProtoLNet is not designed to
 312 outperform all alternative approaches. Instead, it is designed to provide interpretability with a
 313 minimal loss in accuracy.

318 The power of the ProtoLNet is that once trained, its decision-making process can be interpreted
 319 by the user. Three example predictions are shown in Fig. 5, along with their two “most winning”
 320 prototypes (i.e. prototypes that gave the most points to the winning class in each example) and the
 321 associated location scaling grids. To avoid any confusion, we want to clearly state that the “pro-
 322 prototypes” outlined in colored boxes in Fig. 5(i),(iii) are not the prototypes themselves. The actual



314 FIG. 5. Three example predictions by the network for the idealized quadrants use case, along with the two
 315 winning prototypes for each sample and the associated location scaling grid. For each of the three samples, there
 316 are two prototypes shown along with their associated location scaling grids. These are indexed as (i,iii) and
 317 (ii,iv), respectively.

323 prototypes are vectors of latent patches of size $1 \times 1 \times 128$ and would likely be incomprehensible
324 since they capture the output of a series of complex convolutions, poolings, and nonlinear activa-
325 tions. Instead, we visualize the group of neighboring pixels of the training field that contribute
326 to the prototype latent patch, often termed the “receptive field”. In contrast, the location scaling
327 panels in Fig. 5(ii),(iv) display the actual grids used in the prototype layer computation, which is
328 why the squares are much larger than the pixels in the input field (i.e. the dimensions have been
329 reduced to 25×25).

330 Consider Sample 230 (Fig. 5a), which the ProtoLNet correctly labeled as class 0. Prototypes 2
331 and 4 contributed the most points to a class 0 prediction, giving 8.8 and 5.2 points, respectively.
332 Prototype 2 was drawn from training sample 6 and, more specifically, Prototype 2 represents a
333 latent patch from the purple-boxed region of training sample 6 (Fig. 5(ai)). The location scaling
334 grid for Prototype 2 (Fig. 5(aii)) shows that this prototype is highly relevant only when found in
335 the upper-left corner of the field (dark gray and black pixels). Thus, the ProtoLNet identified high
336 similarity between Prototype 2 and an upper-left patch of Sample 230. Or in other words, the
337 ProtoLNet identified that sample 230 *looks like that* prototype *there*.

338 Prototype 4 (Fig. 5(aiii),(iv)) also contributed points to the correct prediction of class 0. Note that
339 Prototype 4 was also drawn from training sample 6; coincidentally the same sample as Prototype
340 2. Looking at Prototypes 2 and 4 together, one can interpret that the network’s decision-making
341 strategy is to look for blue anomalies in the upper-left and bottom-right quadrants of the image
342 — which is exactly how class 0 is defined. A similar interpretation can be found for sample 78
343 (Fig. 5b) with a class label of 1. The network identifies the class 1 sample by looking for positive
344 anomalies in the upper-left and bottom-left quadrants.

345 The network’s decision-making strategy is particularly interesting for Sample 153 with a label of
346 class 2 (Fig. 5c). Prototype 13 corresponds to features associated with a weakly positive anomaly
347 in the upper-left or bottom-right quadrants. From this, it appears that the network is ruling out
348 a class 0 sample, which exhibits negative anomalies in these quadrants. Similarly, Prototype 14
349 corresponds to features associated with a weakly negative anomaly in the upper-left or bottom-left
350 quadrants. That is, the network rules out a class 1 field that exhibits strong positive anomalies
351 in these two quadrants. Fig. 5(cii),(iv) nicely demonstrates that the location scaling grid can
352 highlight multiple locations throughout the field for the same prototype. The interpretability of the

353 ProtoLNet prediction thus allows for identification of the patches of the input field that were used
354 to make the prediction, i.e. the patches whose latent representation most looks like class-specific
355 prototypes learned during training.

356 One interesting observation is the sparsity of the location scaling grids in (Fig. 5) despite no
357 explicit sparsity requirement in the loss function. This comes about due to the SeparationCost (Eq.
358 2) pushing the values of the location scaling grid to lower values in unimportant areas. Since the
359 SeparationCost is subtracted in the loss function, and the location scaling values (s_k) appear in
360 the denominator of the SeparationCost, the gradient of the loss function ultimately favors small
361 location scaling values for unfavorable prototypes.

362 4. Use Case: MJO Phase Classification

363 We next apply the ProtoLNet architecture to earth system reanalysis fields. Specifically, the net-
364 work is tasked with ingesting maps of atmospheric fields in the tropics and predicting the current
365 phase of the Madden-Julian oscillation (MJO). The MJO is a large-scale, eastward propagating cou-
366 pling between tropical wind and convection that oscillates on subseasonal (30-60 day) timescales
367 (Madden and Julian 1971, 1972; Zhang 2005). Canonical MJO events form in the Indian Ocean,
368 and propagate east into the western Pacific: the “phase” of the MJO describes roughly where it is
369 in this life cycle.

370 The task of classifying the current phase of the MJO from maps of the tropics is chosen
371 here to demonstrate the utility of our method to a relatively straightforward climate science task.
372 Classification of MJO phase requires the network to identify coherent, multivariate tropical patterns
373 on a particular (planetary) spatial scale, and the MJO’s eastward propagation also requires the
374 network to take advantage of spatial location in its decision making. Thus, while straightforward
375 from a scientific perspective, the task of classifying MJO phase is well-suited as a demonstrative
376 use-case for the ProtoLNet methodology. Toms et al. (2021) classified the state of the MJO to
377 explore the utility of explainability methods, in contrast to our interpretable method, for earth
378 system science applications.

379 We define MJO activity and phase using the “Real-time Multivariate MJO index” (RMM;
380 Wheeler and Hendon (2004)). RMM is derived through an empirical orthogonal function (EOF)
381 analysis of three variables: outgoing longwave radiation (OLR), 200 hPa zonal wind (u200) and

382 850 hPa zonal wind (u850). Each variable in RMM is pre-processed by removing the seasonal cycle
383 (i.e. the all-time mean and first three harmonics of the annual cycle on each calendar day), and the
384 previous 120-day mean of each day (to remove variability associated with longer timescales than
385 the MJO). Variables are averaged from 15N-15S, and the leading two modes of the EOF analysis
386 are used to define the MJO through two daily time series. Plotted on a 2-dimensional plane, the
387 distance of a point from the origin represents the strength of the MJO (often called the RMM
388 amplitude), and the phase angle describes the phase of the MJO, or where it is in its life cycle.
389 Following Wheeler and Hendon (2004), when the MJO is active (e.g. above a certain amplitude
390 threshold) we divide the RMM phase space into octants. Phases 1 and 2, for example, correspond
391 to active MJO convection in the Indian Ocean. Phases 3 and 4 are associated with activity around
392 the Maritime Continent, etc.. If the MJO is not active, we label it as Phase 0.

393 We define and track the MJO using ERA-20C reanalysis data (Poli et al. 2016), a reanalysis dataset
394 than spans the entire twentieth century and provides a larger sample size than the observational
395 record. From ERA-20C, we use daily OLR, u850, and u200 data from May 1, 1900 until December
396 31, 2010 to calculate the RMM index. RMM is calculated from the ERA-20C data following the
397 methodology in Wheeler and Hendon (2004) discussed above, except that the full ERA-20C period
398 is used to define the climatology, and the processed data are projected onto the *observed* EOF
399 modes from Wheeler and Hendon (2004) (as opposed to the EOFs from the ERA-20C data). Over
400 the period when the observed RMM index overlaps with our ERA-20C RMM index, the two indices
401 have a correlation of approximately 0.9, indicating very good agreement in how the RMM index is
402 formed.

403 The network input is composed of three channels of 17 latitudes by 105 longitudes of u200,
404 u850, and OLR, representing the three geophysical variables that go into the computation of the
405 MJO index (see Fig. 6). Thus, a single sample has shape $17 \times 105 \times 3$. The labels are set to be the
406 phase of the MJO, with phase 0 representing days where the amplitude of the MJO is less than 0.5.
407 We choose to train on all available data; thus, the classes are not equally balanced across phases
408 (see Supp. Fig. S1), although they are similar.

409 Given that there is memory of the MJO phase from one day to the next, we divide the 1900-
410 2010 data into training and testing via distinct years. Specifically, the testing data is all calendar
411 days within the 22 randomly selected years: 1902, 1903, 1907, 1912, 1916, 1917, 1918, 1923,

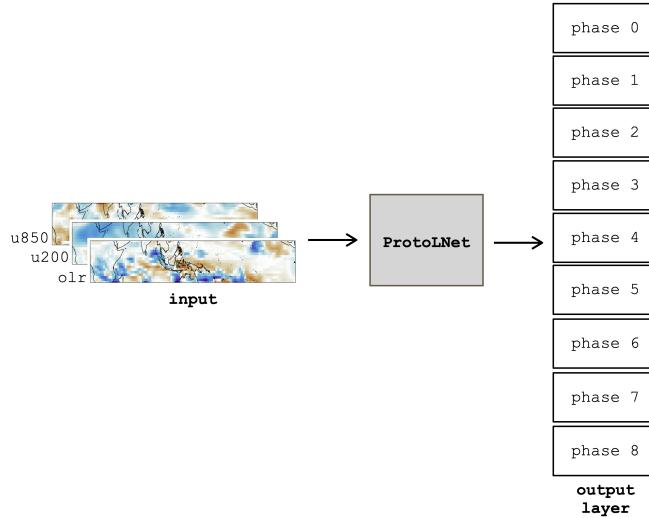


FIG. 6. Prediction setup for the MJO use case.

412 1935, 1937, 1941, 1945, 1946, 1949, 1953, 1961, 1965, 1976, 1992, 2007, 2008, and 2010. The
 413 training years comprise the remaining 89 years. (Results for other combinations of training/testing
 414 accuracies are given in Supp. Table S1.) This results in 32,387 training samples and 8,035 testing
 415 samples. The three input fields (channels) are converted to anomalies prior to analysis following a
 416 similar pre-processing as for the RMM computation. That is, the time-mean calendar-day seasonal
 417 cycle is subtracted from each gridpoint, and the mean of the previous 120 days is removed. Each
 418 variable is individually normalized by dividing it by its tropics-wide standard deviation. Then,
 419 immediately prior to training, the inputs are further standardized by the mean and standard deviation
 420 across all gridpoints and channels of the training set (via flattening the input fields).

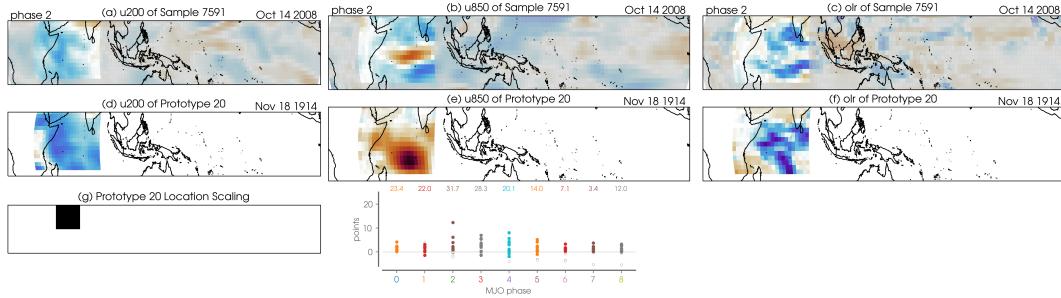
421 We first train a standard CNN to perform the classification task and act as our base CNN for the
 422 ProtoLNet. Details of the CNN architecture and training parameters are provided in Appendix A.
 423 Once the CNN is trained, we remove the final fully connected layer and output layer, and append
 424 the result to the ProtoLNet to become the base CNN (see Fig. 1). We assign 10 prototypes (with
 425 $D = 64$) to each output class, which results in a total of 90 prototypes. Fewer than 90 reduced the
 426 accuracy, while using more than 90 did not improve the predictions. We cycle through the three
 427 stages of ProtoLNet training (Fig. 2) five times, freezing the base CNN for the first cycle of stage
 428 1, but allowing it to train on all subsequent cycles of stage 1. Once fully trained, the ProtoLNet
 429 achieves a testing accuracy of 73% for classifying the phase of the MJO into one of nine classes

430 (random chance is approximately 11%), which is similar to the accuracy found in Toms et al.
431 (2021) using a black box neural network. Supp. Fig. S2 shows that the ProtoLNet exhibits testing
432 accuracies between approximately 70-80% across phases. A ProtoPNet, which does not consider
433 location, never achieves an accuracy above 30%.

434 Interestingly, the base CNN upon which our ProtoLNet was trained converged to an accuracy of
435 58%, much lower than that of the subsequent ProtoLNet. We believe that the improved accuracy of
436 the ProtoLNet is due to the regularizing nature of the prototype architecture. That is, the prototype
437 approach constrains the network to focus on only a few latent features for phase identification,
438 allowing it to converge on an appropriate decision-making strategy when the training data is
439 limited (see discussion of additional experiments in Section 5). We believe that this may be an
440 additional benefit of the prototype approach that is worthy of further investigation. With that said,
441 Supp. Table S1 shows accuracies for the base CNN and ProtoLNet for six additional random seeds
442 that set the model initialization and training/testing split. The ProtoLNet accuracies are incredibly
443 robust across all seeds. While in two of the cases the base CNN achieved lower accuracies than
444 the ProtoLNet (as in the setup shown here), the base CNN more often achieved a slightly higher
445 accuracy than the ProtoLNet. Thus, it appears that the original accuracy of the base CNN does not
446 solely dictate the resulting accuracy of the ProtoLNet.

454 An example of the interpretability of the ProtoLNet’s prediction for testing sample 7591 is shown
455 in Fig. 7. This sample corresponds to phase 2 of the MJO on October 14, 2008, and the three
456 input fields (u200, u850, and olr) are displayed across the top row for that day. All anomalies are
457 shown, but the shading outside of the prototype receptive field is muted in color. Note that the
458 large-scale, enhanced convection of the western Indian Ocean (Fig. 7c) is a classic indication of a
459 phase 2 MJO event, corresponding with a coupled wind response that shows upper-level easterlies
460 (Fig. 7a), and lower-level westerlies (Fig. 7b) in the same region.

461 The network correctly classifies this sample as phase 2, and we can use the interpretability of
462 the ProtoLNet to further explore why. Although multiple prototypes contributed to the winning
463 number of points for the classification of sample 7591, it can be insightful to investigate the winning
464 prototype (i.e. the prototype that contributes the most points). With multiple channels as input,
465 the winning prototype for this sample (Prototype #20) is visualized as three different fields, one
466 for each input variable (i.e. u200, u850, olr), as shown in Fig. 7d-f. Prototype 20 is a latent

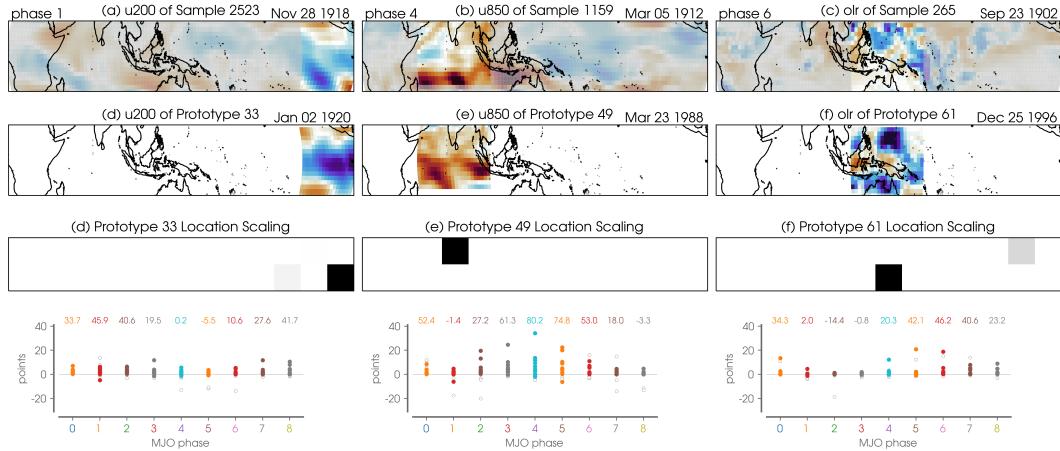


447 FIG. 7. One example prediction (testing sample 7591) by the ProtoLNet for the MJO use case, along with
 448 the winning prototype (prototype 20) and associated location scaling grid. The three columns denote the three
 449 input fields (i.e. u200, u850, olr). All anomalies are shown in panels (a)-(c) but the shading outside of the
 450 prototype patch is muted in color. The color scales are dimensionless with red shading denoting positive values
 451 and blue shading denoting negative values. The bottom middle panel show the points given to each class by each
 452 prototype, with the sum (i.e. total points) displayed along the top. Colored dots denote prototypes associated
 453 with the same class, and white dots denote contributions from prototypes of other classes.

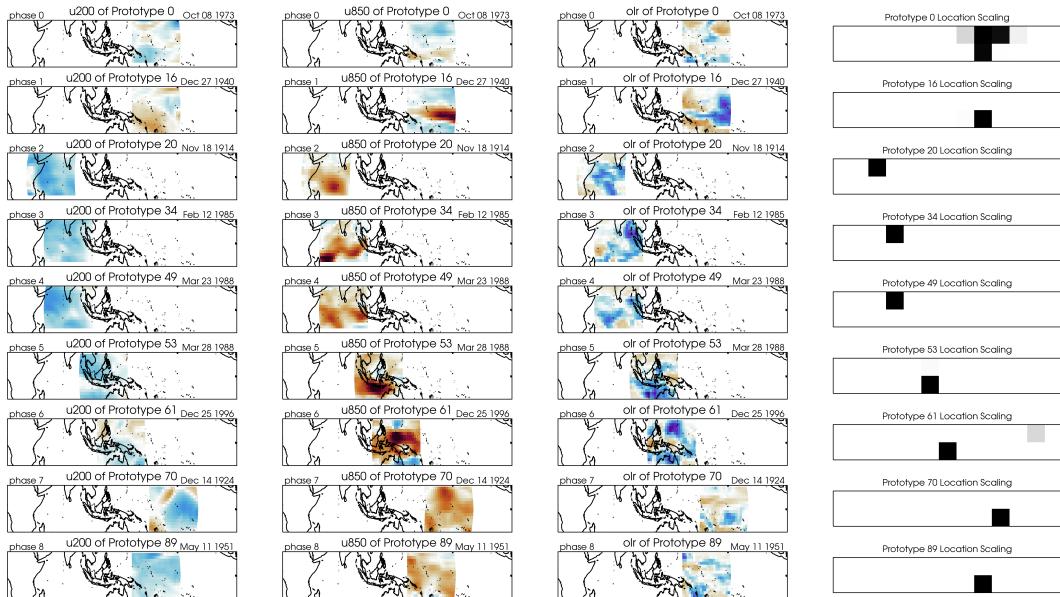
467 patch corresponding to the state of the western Indian Ocean on November 18, 1914. The location
 468 scaling grid associated with Prototype 20 (Fig. 7g) highlights that similarities to this prototype are
 469 only heavily weighted when found at these longitudes. Thus, we see that the anomaly fields on
 470 October 14, 2008, for sample 7591 look a lot like those of Prototype 20, with upper-level easterlies,
 471 lower-level westerlies and enhanced convection over the western Indian Ocean. This provides
 472 evidence for why the network classified this sample as MJO phase 2.

475 Fig. 8 shows three additional (correctly predicted) testing samples and their winning prototypes,
 476 displaying only one geophysical field for each prediction to simplify the figure. Sample 2523 on
 477 November 28, 1918, is classified as phase 1, in part because its upper-level easterlies *look like*
 478 those of Prototype 33 from January 2, 1920 over the eastern Pacific (Fig. 8a,d). The lower-level
 479 westerlies over the Indian Ocean on March 5, 1912, *look like* those of phase 4 Prototype 49 from
 480 March 23, 1988 (Fig. 8b,e). Enhanced convection as seen by the OLR field east of the Maritime
 481 Continent on September 23, 1902 *looks like* that of phase 6 Prototype 61 (Fig. 8c,f).

484 As a summary of the MJO classification results, Fig. 9 displays the most frequently winning
 485 prototype for each phase of the MJO. A hallmark feature of the MJO is its eastward propagation,
 486 and Fig. 9 reveals the eastward progression of the prototypes (and associated location scaling

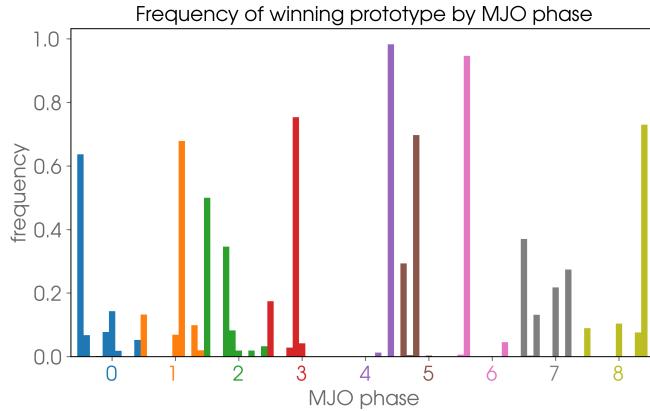


473 FIG. 8. As in Fig. 7, but for three additional example testing samples (one per column) displaying only one
 474 geophysical field for each.



482 FIG. 9. The most frequently winning prototype for correctly classified testing samples by MJO phase. Each
 483 column represents a different input variable; the fourth column displays the associated location scaling.

487 grids) starting in phase 2 and continuing to phases 7 and 8. That is, the ProtoLNet, with its
 488 location-specific focus, has learned representative prototypes that move eastward with the known
 489 progression of the MJO. Phase 1, however, does not appear to behave this way. Prototype 16 is
 490 often the most-winning prototype for phase 1, but it is focused over the mid-Pacific rather than

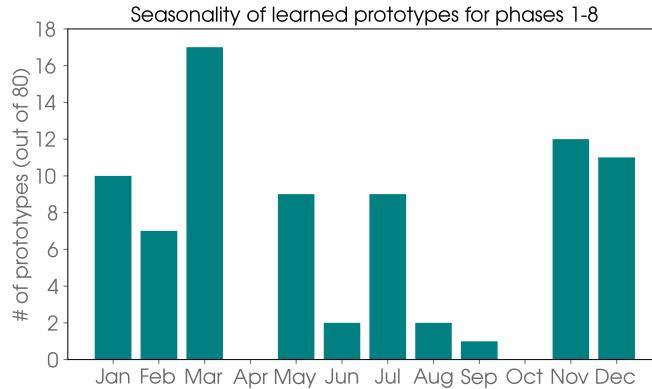


495 FIG. 10. The frequency that each prototype is the winning prototype (i.e. contributes the most points to the
 496 predicted class) for each correctly classified testing sample. Each phase has 10 possible prototypes; however,
 497 there are some prototypes that are never a winning prototype. They have frequency of zero.

491 the western Indian Ocean as one might expect (this is true for most of the Phase 1 prototypes; see
 492 Supp. Fig. S5). The reason why phase 1 prototypes tend to focus on this region is not clear, but we
 493 hypothesize the network may be focusing on wind signals in this region associated with a phase 1
 494 event forming or a previous MJO event decaying. Further investigation is needed.

498 Fig. 10 shows a breakdown of how often (i.e. for how many testing samples) each prototype
 499 was the winning prototype. For example, Prototype 49 from March 23, 1988, is the most-winning
 500 prototype for phase 4, and it is the winning prototype for 98% of all correctly classified phase 4
 501 testing samples. This suggests that this prototype is highly indicative of phase 4 MJO events. On
 502 the other hand, phase 7 has multiple prototypes that frequently earn the title of winning prototype.
 503 Thus, Prototype 70 (displayed in Fig. 9) should be interpreted as only one possible indicator of
 504 phase 7.

505 All 10 learned prototypes for each phase are provided in the Supp. Fig. S4-S12. Careful
 506 inspection shows that some of the learned prototypes come from the same training sample, indi-
 507 cating a particularly prototypical event. However, in cases where the prototypes come from the
 508 same training sample and have similar location scaling grids, there could be concern that this is
 509 a repeated prototype. Chen et al. (2019) discuss an additional “pruning” step in their ProtoPNet
 510 methodology, although it could also be that the CNN is identifying different aspects of the image
 511 that are prototypical. Either way, for this MJO use case the Supp. Fig. S4-S12 specify how often



514 FIG. 11. Number of learned prototypes for MJO phases 1-8 (excluding phase 0, so out of 80 prototypes total)
 515 binned by month of the year of the training sample from which the prototype was drawn.

512 a particular prototype was the “winning” prototype and in all cases it is one specific prototype that
 513 wins-out over the rest which is why we are confident showing Fig. 10.

516 Fig. 11 shows the breakdown of the monthly distribution for all prototypes for active MJO phases
 517 1-8. The network preferentially chooses prototypes from November-March when the MJO is
 518 known to be most active, however, prototypes from May and July are also learned, likely to capture
 519 the differences in MJO behavior across seasons (Zhang 2005). The monthly seasonality for all
 520 prototypes, including those for MJO phase 0, are shown in Supp. Fig. S3.

521 5. Discussion

522 The value of the ProtoLNet design is that interpretation of the network’s decision-making process
 523 is baked into the architecture itself, rather than performed *post-hoc* like most explainable AI methods
 524 (Buhrmester et al. 2019; Barredo Arrieta et al. 2020; Samek et al. 2021). Although the network
 525 is constrained to only learn via similarities to a small number of learned prototypes, multiple use
 526 cases demonstrate that it can be trained to exhibit only a small reduction in accuracy compared to
 527 non-interpretable architectures (Chen et al. 2019; Singh and Yow 2021). Moreover, for our MJO
 528 use case, the ProtoLNet actually improved in accuracy over its base CNN. We hypothesize that this
 529 is because the ProtoLNet greatly reduces the search space possibilities, which allows the network
 530 to converge on a good prediction strategy given a limited sample size. One might think of this as
 531 a form of regularization, or instead, a form of physics-guided constraint (e.g. Beucler et al. 2021)

532 that forces the network to learn physically realizable evidence for each class. To further explore this
533 hypothesis, we trained additional ProtoLNets for the idealized quadrants use case (Section 3), but
534 with a much smaller training size (only 1,400 samples for training). In all cases, the ProtoLNets
535 obtained higher testing accuracies — sometimes significantly higher — than their respective base
536 CNNs (see results in Supp. Fig. S13). This is not to say that the ProtoLNet is categorically more
537 accurate than a standard CNN. A more thorough exploration of the hyperparameter space could
538 bring the base CNN accuracy up to that of the ProtoLNet. Instead, we just wish to highlight that
539 with minimal tuning, the ProtoLNet was able to consistently achieve high accuracies with limited
540 training data.

541 In addition to being interpretable, the ProtoLNet provides the benefit of learning a small subset
542 of prototypical parts *from the training set* that reflect identifiable features for each output class.
543 That is, each prototype is found “in the wild” and, thus, has a direct connection to a sample that
544 has occurred. This should be distinguished from more standard architectures that learn complex
545 latent representations and features that may never occur in reality. For the case of MJO phase
546 classification, this means that the network can learn particular example MJO events that generalize
547 across the observational record and reflect identifiable features for each specific MJO phase.
548 Thus, although predicting the current phase of the MJO is routine from a scientific perspective,
549 the ProtoLNet allows us to look back and identify specific dates that exhibit prototypical MJO
550 phase behaviour, as shown in Fig. 9 and Fig. 11. Furthermore, it is straightforward to extend the
551 interpretable ProtoLNet setup of Fig. 6 to ingest current atmospheric fields and predict the MJO
552 phase at some lead time into the future.

553 As we have used it here, the ProtoLNet design learns localized prototypes from the input that
554 provide evidence for a particular output class. This should be distinguished from the standard
555 climate approach that composites the input fields over many samples for a single class, and thus
556 results in a smooth averaged field (assuming there are enough samples to average out the noise).
557 Such a composite field is computed pixel by pixel and as such, does not capture shared gradients or
558 higher-level features that can be learned by the convolutional layers of the ProtoLNet. Finally, as
559 discussed above, the ProtoLNet identifies prototypical behavior that has been realized in a training
560 sample, while the composite field provides a smoothed, idealized picture that will likely never be
561 observed.

562 The ProtoLNet is based on the ProtoPNet of Chen et al. (2019) which uses positive reasoning,
563 i.e. *this looks like that*, to predict the correct class of an input image. Singh and Yow (2021)
564 introduce a variation, the NP-ProtoPNet, which additionally includes negative reasoning, i.e. *this*
565 *does not look like that*. Their argument is that by allowing negative reasoning, the network is able
566 to better rule out incorrect classes and achieve accuracies on-par with the best performing black box
567 models. It is straightforward to apply our location-scaling grid to a NP-ProtoPNet, which mainly
568 involves relaxing the sparsity requirement of the final weights layer. However, by allowing both
569 positive and negative reasoning, interpreting the model’s decision making process may become
570 significantly more difficult due to competing negative and positive point contributions to the final
571 output classes. Thus, we chose to focus on positive reasoning for this study.

572 **6. Conclusions**

573 Driven by the desire to explain the decision-making process of deep learning models, a large
574 variety of *post hoc* explainability methods have been developed (e.g. Buhrmester et al. 2019;
575 Barredo Arrieta et al. 2020; Samek et al. 2021). However, these explainability methods come with
576 their own challenges (Kindermans et al. 2019; Mamalakis et al. 2021) and recent work by Rudin
577 (2019) and Chen et al. (2019) suggest that instead of trying to explain black box models, we should
578 be creating models where the decision-making process is interpretable by design.

579 Here, we extend the interpretable ProtoPNet of Chen et al. (2019) to consider absolute location
580 in the interpretable prototype architecture, which we term the ProtoLNet. The results of our
581 work can be summarized by three main conclusions. (1) Considering absolute location in the
582 ProtoLNet architecture greatly improves accuracy for the geophysical use cases explored here. (2)
583 The ProtoLNet is interpretable in that it directly provides which prototypes are similar to different
584 patches of an input image (i.e. *this looks like that*), and where these prototypes matter (i.e. *there*).
585 (3) The network is able to learn specific historical dates that serve as multivariate prototypes of the
586 different Madden-Julian oscillation phases.

587 This work serves as one example of an interpretable deep learning model specifically designed
588 for earth system science applications (see also Sonnewald and Lguensat 2021). There is much
589 more research to be done on the topic. For example, the incorporation of negative reasoning
590 and extension to regression tasks could be beneficial for its use in earth science. Furthermore,

591 the interpretation and utility of the learned prototypes themselves, apart from the prediction task,
592 leaves much to be explored. Thus, this work should be seen as merely a step in the direction of
593 interpretable deep learning for earth science exploration.

594 *Acknowledgments.* This work was funded, in part, by the NSF AI Institute for Research on
595 Trustworthy AI in Weather, Climate, and Coastal Oceanography (AI2ES) under NSF grant ICER-
596 2019758. ZKM recognizes support from the National Science Foundation under Award No.
597 2020305. JKR recognizes support from the U.S. Department of Energy, Office of Science, Office of
598 Advanced Scientific Computing Research, Department of Energy Computational Science Graduate
599 Fellowship under Award No. DE-SC0020347.

600 *Data availability statement.* Once published, the code will be made available to the community
601 via a permanent DOI on Zenodo. For peer-review, the code is available at [https://github.com/
602 eabarnes1010/tlltt](https://github.com/eabarnes1010/tlltt). The ERA-20C is publically available at [https://www.ecmwf.int/en/
603 forecasts/datasets/reanalysis-datasets/era-20c](https://www.ecmwf.int/en/forecasts/datasets/reanalysis-datasets/era-20c).

604 APPENDIX A

605 **Base CNN architectures and training**

606 The base CNN for the idealized quadrants use case (Section 3) has two convolutional layers of 32
607 kernels each. Every convolutional layer is followed by an average pooling layer with kernel size
608 2×2 and a stride length of 2. The output of the final average pooling layer is flattened, and then
609 fed into a final dense layer of 64 units which is fed into the final output layer of 3 units. The final
610 output layer contains the softmax activation function which convert the outputs into confidences
611 that sum to 1.0. The final dense layer is trained with dropout (Srivastava et al. 2014) at a rate of
612 0.4 to reduce overfitting. When the base CNN is appended to the ProtoLNet, the dropout rate is
613 set to zero. That is, dropout is only used to reduce overfitting during the pre-training of the base
614 CNN. The base CNN is trained with a fixed learning rate of $5e-5$ for 12 epochs.

615 The base CNN for the MJO use case (Section 4) has three convolutional layers of 16 kernels
616 each. Every convolutional layer is followed by an average pooling layer with kernel size 2×2 and
617 a stride length of 2. The convolutional layers are trained with dropout at a rate of 0.4 to reduce
618 overfitting. The output of the final average pooling layer is flattened, and then fed into a final dense
619 layer of 32 units that is fed into the final output layer of 9 units. The final output layer contains
620 the softmax activation function which converts the outputs into confidences that sum to 1.0. The
621 final dense layer is trained with dropout at a rate of 0.2. When the base CNN is appended to the
622 ProtoLNet, the dropout rates are set to zero. That is, dropout is only used to reduce overfitting

623 during the pre-training of the base CNN. The base CNN is trained with a fixed learning rate of
 624 0.00017548 for 23 epochs.

625 The hyperparameters for these networks were explored using KerasTuner. We did not find the
 626 results to be overly sensitive to these choices.

627 APPENDIX B

628 Learning location scaling exponents

629 The location scaling values must be non-negative. Subsequently, we use a trick from Duerr et al.
 630 (2020) and learn the exponents of the location scaling, rather than the values themselves. That is,
 631 if s_k denotes the location scaling value for prototype \mathbf{p} at latent patch k then

$$632 \quad s_k = e^{\gamma_k}, \quad (\text{B1})$$

633 where the free parameter γ_k is *learned* by the network during training. Thus, at initialization, all
 634 γ_k values are initialized to zero so that the location scaling grid (all s_k values) is initialized to a
 635 grid of ones.

636 APPENDIX C

637 Stage 1 loss function

638 The Stage 1 loss function is given by Equation 2. There are three components: the usual CrossEn-
 639 tropy, plus a ClusterCost, and minus a SeparationCost.

640 Consider a set of input samples and associated class labels $\{(\mathbf{x}_i, y_i) : i = 1, 2, \dots, N\}$. The output
 641 from the extended CNN given sample \mathbf{x}_i is a quilt of latent patches \mathbf{z}_{ik} , where k indexes the latent
 642 patches. For the architecture shown in Figure 1, $k \in \{1, 2, \dots, 6\}$ because the quilt is 2×3 . Let s_k
 643 denote the current location scaling value associated with latent patch k , and \mathbf{P}_{y_i} denote the set of
 644 all prototypes belonging to class y_i . The ClusterCost is given by

$$645 \quad \text{ClusterCost} = \frac{1}{N} \sum_{i=1}^N \left[\min_{\mathbf{p} \in \mathbf{P}_{y_i}} \min_k \frac{\|\mathbf{z}_{ik} - \mathbf{p}\|_2^2}{s_k + \epsilon} \right] \quad (\text{C1})$$

646 where $\|\cdot\|_2^2$ is the squared L_2 norm and ϵ is a small number, there to guard against divide-by-zero
647 problems.

648 The ClusterCost encourages training images to have at least one latent patch with high similarity
649 to a prototype of the same class. The computation is based on Chen et al. (2019), but incorporates
650 the location scaling grid introduced in this paper.

651 The SeparationCost discourages training images from having high similarity to prototypes be-
652 longing to the incorrect class. The computation is almost identical to that of the ClusterCost. The
653 difference is that we minimize over the set of all prototypes that do not belong to class y_i .

$$654 \text{SeparationCost} = \frac{1}{N} \sum_{i=1}^N \left[\min_{\mathbf{p} \notin \mathbf{P}_{y_i}} \min_k \frac{\|\mathbf{z}_{ik} - \mathbf{p}\|_2^2}{s_k + \epsilon} \right] \quad (\text{C2})$$

655 For the idealized quadrants use case, we set the ClusterCost coefficient $\beta_1 \approx 0.17$ (see code for
656 all digits) and the SeparationCost coefficient $\beta_2 = \beta_1/10$. For the MJO use case $\beta_1 = 0.2$ and
657 $\beta_2 = \beta_1/10$. Note the negative sign in front of the SeparationCost term in Equation 2 encourages
658 the network to have larger separation (lower similarity) between samples and the prototypes from
659 incorrect classes.

660 References

661 Balmaseda, M., and Coauthors, 2020: NOAA-DOE precipitation processes and predictability
662 workshop. Tech. Rep. DOE/SC-0203; NOAA Technical Report OAR CPO-9., U.S. Department
663 of Energy and U.S. Department of Commerce NOAA.

664 Barnes, E. A., B. Toms, J. W. Hurrell, I. Ebert-Uphoff, C. Anderson, and D. Anderson, 2020:
665 Indicator Patterns of Forced Change Learned by an Artificial Neural Network. *Journal of*
666 *Advances in Modeling Earth Systems*, **n/a (n/a)**, e2020MS002195, <https://doi.org/10.1029/2020MS002195>.

668 Barredo Arrieta, A., and Coauthors, 2020: Explainable artificial intelligence (XAI): Concepts,
669 taxonomies, opportunities and challenges toward responsible AI. *Inf. Fusion*, **58**, 82–115,
670 <https://doi.org/10.1016/j.inffus.2019.12.012>.

- 671 Beucler, T., M. Pritchard, S. Rasp, J. Ott, P. Baldi, and P. Gentine, 2021: Enforcing analytic
672 constraints in neural networks emulating physical systems. *Phys. Rev. Lett.*, **126** (9), 098 302,
673 <https://doi.org/10.1103/PhysRevLett.126.098302>.
- 674 Buhrmester, V., D. Münch, and M. Arens, 2019: Analysis of explainers of black box deep neural
675 networks for computer vision: A survey. *arXiv*, 1911.12116.
- 676 Chen, C., O. Li, D. Tao, A. Barnett, C. Rudin, and J. K. Su, 2019: This looks like that: Deep learning
677 for interpretable image recognition. *Advances in Neural Information Processing Systems*, Curran
678 Associates, Inc., Vol. 32.
- 679 Davenport, F. V., and N. S. Duffenbaugh, 2021: Using machine learning to analyze physical causes
680 of climate change: A case study of U.S. midwest extreme precipitation. *Geophys. Res. Lett.*,
681 <https://doi.org/10.1029/2021gl093787>.
- 682 Duerr, O., B. Sick, and E. Murina, 2020: *Probabilistic Deep Learning: With Python, Keras and*
683 *Tensorflow Probability*. MANNING PUBN.
- 684 Géron, A., 2019: *Hands-on Machine Learning with Scikit-Learn, Keras, and TensorFlow*. O'Reilly
685 UK Ltd.
- 686 He, K., X. Zhang, S. Ren, and J. Sun, 2015: Delving deep into rectifiers: Surpassing Human-Level
687 performance on ImageNet classification. *2015 IEEE International Conference on Computer*
688 *Vision (ICCV)*, 1026–1034, <https://doi.org/10.1109/ICCV.2015.123>.
- 689 Irrgang, C., N. Boers, M. Sonnewald, E. A. Barnes, C. Kadow, J. Staneva, and J. Saynisch-
690 Wagner, 2021: Towards neural earth system modelling by integrating artificial intelligence in
691 earth system science. *Nature Machine Intelligence*, **3** (8), 667–674, <https://doi.org/10.1038/s42256-021-00374-3>.
- 693 Keys, P. W., E. A. Barnes, and N. H. Carter, 2021: A machine-learning approach to human
694 footprint index estimation with applications to sustainable development. *Environ. Res. Lett.*,
695 **16** (4), 044 061, <https://doi.org/10.1088/1748-9326/abe00a>.
- 696 Kindermans, P.-J., S. Hooker, J. Adebayo, M. Alber, K. T. Schütt, S. Dähne, D. Erhan, and
697 B. Kim, 2019: The (un)reliability of saliency methods. *Explainable AI: Interpreting, Explaining*

698 *and Visualizing Deep Learning*, W. Samek, G. Montavon, A. Vedaldi, L. K. Hansen, and K.-
699 R. Muller, Eds., Springer International Publishing, Cham, 267–280, [https://doi.org/10.1007/](https://doi.org/10.1007/978-3-030-28954-6_14)
700 [978-3-030-28954-6_14](https://doi.org/10.1007/978-3-030-28954-6_14).

701 Lapuschkin, S., S. Wäldchen, A. Binder, G. Montavon, W. Samek, and K.-R. Müller, 2019:
702 Unmasking clever hans predictors and assessing what machines really learn. *Nat. Commun.*,
703 **10 (1)**, 1096, <https://doi.org/10.1038/s41467-019-08987-4>.

704 Madden, R. A., and P. R. Julian, 1971: Detection of a 40-50 day oscillation in the zonal wind
705 in the tropical pacific. *Journal of Atmospheric Sciences*, **28 (5)**, 702 – 708, [https://doi.org/](https://doi.org/10.1175/1520-0469(1971)028<0702:DOADOI>2.0.CO;2)
706 [10.1175/1520-0469\(1971\)028<0702:DOADOI>2.0.CO;2](https://doi.org/10.1175/1520-0469(1971)028<0702:DOADOI>2.0.CO;2), URL [https://journals.ametsoc.org/](https://journals.ametsoc.org/view/journals/atsc/28/5/1520-0469_1971_028_0702_doadoi_2_0_co_2.xml)
707 [view/journals/atsc/28/5/1520-0469_1971_028_0702_doadoi_2_0_co_2.xml](https://journals.ametsoc.org/view/journals/atsc/28/5/1520-0469_1971_028_0702_doadoi_2_0_co_2.xml).

708 Madden, R. A., and P. R. Julian, 1972: Description of global-scale circulation cells in the tropics
709 with a 40-50 day period. *Journal of Atmospheric Sciences*, **29 (6)**, 1109 – 1123, [https://doi.org/](https://doi.org/10.1175/1520-0469(1972)029<1109:DOGSCC>2.0.CO;2)
710 [10.1175/1520-0469\(1972\)029<1109:DOGSCC>2.0.CO;2](https://doi.org/10.1175/1520-0469(1972)029<1109:DOGSCC>2.0.CO;2), URL [https://journals.ametsoc.org/](https://journals.ametsoc.org/view/journals/atsc/29/6/1520-0469_1972_029_1109_dogsc_2_0_co_2.xml)
711 [view/journals/atsc/29/6/1520-0469_1972_029_1109_dogsc_2_0_co_2.xml](https://journals.ametsoc.org/view/journals/atsc/29/6/1520-0469_1972_029_1109_dogsc_2_0_co_2.xml).

712 Mamalakis, A., E. A. Barnes, and I. Ebert-Uphoff, 2022: Investigating the fidelity of explainable
713 artificial intelligence methods for applications of convolutional neural networks in geoscience.
714 2202.03407.

715 Mamalakis, A., I. Ebert-Uphoff, and E. A. Barnes, 2021: Neural Network Attribution Methods
716 for Problems in Geoscience: A Novel Synthetic Benchmark Dataset. *arXiv*, 2103.10005, 2103.
717 10005.

718 Martin, Z. K., E. A. Barnes, and E. D. Maloney, 2021: Using simple, explainable neural networks to
719 predict the madden-julian oscillation. *Earth and Space Science Open Archive*, 42, [https://doi.org/](https://doi.org/10.1002/essoar.10507439.2)
720 [10.1002/essoar.10507439.2](https://doi.org/10.1002/essoar.10507439.2), URL <https://doi.org/10.1002/essoar.10507439.2>.

721 Mayer, K. J., and E. A. Barnes, 2021: Subseasonal forecasts of opportunity identified by an
722 explainable neural network. *Geophys. Res. Lett.*, <https://doi.org/10.1029/2020gl092092>.

723 McGovern, A., R. Lagerquist, D. John Gagne, G. E. Jergensen, K. L. Elmore, C. R. Homeyer,
724 and T. Smith, 2019: Making the black box more transparent: Understanding the physical

725 implications of machine learning. *Bull. Am. Meteorol. Soc.*, **100** (11), 2175–2199, [https://doi.org/](https://doi.org/10.1175/BAMS-D-18-0195.1)
726 10.1175/BAMS-D-18-0195.1.

727 Montavon, G., W. Samek, and K.-R. Müller, 2018: Methods for interpreting and understanding deep
728 neural networks. *Digit. Signal Process.*, **73**, 1–15, <https://doi.org/10.1016/j.dsp.2017.10.011>.

729 National Academies of Sciences Engineering and Medicine, 2020: *Earth System Predictability*
730 *Research and Development: Proceedings of a Workshop–in Brief*. The National Academies
731 Press, Washington, DC, <https://doi.org/10.17226/25861>.

732 Philander, S. G. H., 1983: El niño southern oscillation phenomena. *Nature*, **302** (5906), 295–301,
733 <https://doi.org/10.1038/302295a0>.

734 Poli, P., and Coauthors, 2016: Era-20c: An atmospheric reanalysis of the twentieth century.
735 *Journal of Climate*, **29** (11), 4083 – 4097, <https://doi.org/10.1175/JCLI-D-15-0556.1>, URL
736 <https://journals.ametsoc.org/view/journals/clim/29/11/jcli-d-15-0556.1.xml>.

737 Rasp, S., H. Schulz, S. Bony, and B. Stevens, 2019: Combining crowd-sourcing and deep learning
738 to understand meso-scale organization of shallow convection. *arXiv*, 1906.01906.

739 Rudin, C., 2019: Stop explaining black box machine learning models for high stakes decisions and
740 use interpretable models instead. *Nature Machine Intelligence*, **1** (5), 206–215, [https://doi.org/](https://doi.org/10.1038/s42256-019-0048-x)
741 10.1038/s42256-019-0048-x.

742 Samek, W., G. Montavon, S. Lapuschkin, C. J. Anders, and K.-R. Müller, 2021: Explaining
743 deep neural networks and beyond: A review of methods and applications. *Proc. IEEE*, **109** (3),
744 247–278, <https://doi.org/10.1109/JPROC.2021.3060483>.

745 Simonyan, K., and A. Zisserman, 2014: Very deep convolutional networks for Large-Scale image
746 recognition. *arXiv*, 1409.1556.

747 Singh, G., and K.-C. Yow, 2021: These do not look like those: An interpretable deep learning
748 model for image recognition. *IEEE Access*, **9**, 41 482–41 493, [https://doi.org/10.1109/ACCESS.](https://doi.org/10.1109/ACCESS.2021.3064838)
749 2021.3064838.

- 750 Sonnewald, M., and R. Lguensat, 2021: Revealing the impact of global heating on north atlantic
751 circulation using transparent machine learning. *J. Adv. Model. Earth Syst.*, **13** (8), [https://doi.org/](https://doi.org/10.1029/2021ms002496)
752 10.1029/2021ms002496.
- 753 Srivastava, N., G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, 2014: Dropout: a
754 simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, **15** (1), 1929–1958.
- 755 Toms, B. A., E. A. Barnes, and I. Ebert-Uphoff, 2020: Physically Interpretable Neural Networks
756 for the Geosciences: Applications to Earth System Variability. *Journal of Advances in Modeling*
757 *Earth Systems*, <https://doi.org/10.1029/2019MS002002>.
- 758 Toms, B. A., K. Kashinath, D. Yang, and Prabhat, 2021: Testing the reliability of interpretable
759 neural networks in geoscience using the Madden–Julian oscillation. *Geosci. Model Dev.*, **14** (7),
760 4495–4508, <https://doi.org/10.5194/gmd-14-4495-2021>.
- 761 Wheeler, M. C., and H. H. Hendon, 2004: An all-season real-time multivariate mjo in-
762 dex: Development of an index for monitoring and prediction. *Monthly Weather Review*,
763 **132** (8), 1917 – 1932, [https://doi.org/10.1175/1520-0493\(2004\)132<1917:AARMMI>2.0.CO;](https://doi.org/10.1175/1520-0493(2004)132<1917:AARMMI>2.0.CO;2)
764 2, URL [https://journals.ametsoc.org/view/journals/mwre/132/8/1520-0493_2004_132_1917_](https://journals.ametsoc.org/view/journals/mwre/132/8/1520-0493_2004_132_1917_aarmmi_2.0.co_2.xml)
765 [aarmmi_2.0.co_2.xml](https://journals.ametsoc.org/view/journals/mwre/132/8/1520-0493_2004_132_1917_aarmmi_2.0.co_2.xml).
- 766 Zhang, C., 2005: Madden-julian oscillation. *Reviews of Geophysics*, **43** (2),
767 <https://doi.org/https://doi.org/10.1029/2004RG000158>, URL [https://agupubs.onlinelibrary.](https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2004RG000158)
768 [wiley.com/doi/abs/10.1029/2004RG000158](https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2004RG000158), [https://agupubs.onlinelibrary.wiley.com/doi/pdf/](https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2004RG000158)
769 [10.1029/2004RG000158](https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2004RG000158).