

Subseasonal Forecasts of Opportunity Identified by an Explainable Neural Network

Kirsten J. Mayer¹ and Elizabeth A. Barnes¹

¹Department of Atmospheric Science, Colorado State University, Fort Collins, CO, USA

Key Points:

- Neural networks can be used to identify forecasts of opportunity for subseasonal prediction
- Neural network explainability techniques pinpoint relevant tropical regions for predictions in the North Atlantic
- Clustering of neural network relevance heat maps reveals a potential new forecast of opportunity for the North Atlantic

Corresponding author: Kirsten J. Mayer, kjmayer@rams.colostate.edu

Abstract

Midlatitude prediction on subseasonal timescales is difficult due to the chaotic nature of the atmosphere and often requires the identification of favorable atmospheric conditions that may lead to enhanced skill (“forecasts of opportunity”). Here, we demonstrate that an artificial neural network can identify such opportunities for tropical-extratropical circulation teleconnections within the North Atlantic (40°N, 325°E) at a lead of 22 days using the network’s confidence in a given prediction. Furthermore, layer-wise relevance propagation, an ANN explainability technique, pinpoints the relevant tropical features the ANN uses to make accurate predictions. We find that layer-wise relevance propagation identifies tropical hot spots that correspond to known favorable regions for midlatitude teleconnections and reveals a potential new pattern for prediction in the North Atlantic on subseasonal timescales.

Plain Language Summary

Weather forecasting on 2 week to 2 month timescales is known for its lack of predictability due to the chaotic nature of the atmosphere. One way to improve prediction skill on these timescales involves the identification of periods of atmospheric conditions that lead to enhanced predictability (“forecasts of opportunities”). Here, we show that a neural network can accurately identify these opportunities when trying to predict the atmospheric circulation over the North Atlantic Ocean 4 weeks in advance. A neural network explainability technique is then used to uncover what the network has “learned” to make these accurate predictions. We show that the network identifies known patterns of storminess ideal for midlatitude prediction and uncovers a possible new favorable region for enhanced prediction.

1 Introduction

Subseasonal timescales (2 weeks - 2 months) are known for their lack of predictability (Mariotti et al., 2018), yet reliable and actionable information on these timescales are required for decision making in many sectors such as public health and water management (e.g. Vitart et al., 2012; White et al., 2017). Over the past decade, there has been a substantial research effort to improve prediction on these timescales (e.g. Vitart et al., 2012; Robertson et al., 2015; Vitart et al., 2017; Pegion et al., 2019). One area of subseasonal prediction research focuses on forecasts of opportunity, the idea that certain earth system conditions provide opportunities for enhanced subseasonal prediction skill (Mariotti et al. 2020). When these opportunities arise, the information provided by the earth system’s state can then be leveraged to improve forecast skill. For example, when the Madden-Julian Oscillation (MJO; Madden and Julian (1971, 1972)), a propagating tropical convective phenomenon, is active, its convective heating can lead to the excitation of quasi-stationary Rossby waves (Hoskins and Ambrizzi 1993) that subsequently modulate the midlatitude circulation over the first few weeks following MJO activity (e.g., Hoskins and Karoly, 1981; Sardeshmukh and Hoskins, 1988; Henderson et al., 2016; Baggett et al., 2017; Zheng et al., 2018). When opposing convective anomalies are located over the Indian Ocean and western Pacific (defined as phases 2, 3, 6, and 7), the MJO has been shown to lead to more coherent and consistent modulations of midlatitude weather on subseasonal timescales and consequently, enhanced prediction skill (Tseng et al., 2018). Using the strength and location of tropical convective activity of the MJO to identify periods of enhanced midlatitude prediction skill is, therefore, an example of forecast of opportunity identification. Mundhenk et al. (2018) also show that an empirical model, which solely uses information about the state of the MJO and the Quasi-Biennial Oscillation, outperforms a state-of-the-art numerical prediction model for prediction of atmospheric river activity on subseasonal timescales. This highlights the importance of statistical models for enhancing subseasonal prediction.

Albers and Newman (2019) demonstrate a technique for forecast of opportunity identification through the utilization of expected skill from a linear inverse model. The study demonstrates the ability of the linear statistical model to identify forecasts of opportunity, and raises the question of whether other statistical models, such as artificial neural networks (ANNs), can identify forecasts of opportunity for subseasonal prediction. ANNs are very good at nonlinear function estimation (Chen & Chen, 1995), and thus, may be able to identify both linear and nonlinear relationships that lend predictability. Recently, ANNs have been successfully applied to seasonal prediction of meteorological variables such as monthly rainfall (Abbot & Marohasy, 2014) and surface temperature (Toms et al., 2020) as well as yearly prediction of the El Niño Southern Oscillation (Ham et al., 2019), suggesting ANNs may be useful for identifying subseasonal forecasts of opportunity as well.

In this paper, we test whether an ANN can be used for subseasonal forecast of opportunity identification. To do so, we input tropical outgoing longwave radiation (OLR) anomalies into an ANN and task the network to predict the sign of 500 hPa geopotential height (z500) anomalies in the North Atlantic (40°N, 325°E) 22 days later (e.g. Week 4). Tropical OLR is used to explore the ability of an ANN to identify known relationships between the MJO and the North Atlantic via tropical-extratropical teleconnections (e.g. Cassou, 2008; Henderson et al., 2016). We demonstrate that an ANN can identify subseasonal forecasts of opportunity related to tropical OLR, and through an ANN explainability technique, demonstrate that the ANN identifies these known MJO-like OLR patterns. In addition, we find a possible new tropical OLR pattern associated with predictable behavior of the North Atlantic circulation on subseasonal timescales.

2 Data and Methods

2.1 Data

We use daily mean OLR (1979-2019) from the National Center for Atmospheric Research/National Oceanic and Atmospheric Administration (NCAR/NOAA; Liebmann and Smith (1996)) and daily mean z500 (1979-2019) from the European Centre for Medium-Range Weather Forecasts (ECMWF) Interim reanalysis (ERA-I; Dee et al. (2011)). MJO teleconnections tend to be stronger during boreal winter (Madden, 1986), and therefore, the extended boreal winter months (November-February) are used for the OLR fields. Since we task the network to predict the sign of the z500 anomaly 22 days following a given OLR field, March is also included in the z500 analysis (see Text S1 for reasoning behind the choice of lead).

The annual cycle is removed from both the z500 and OLR data. For z500, the annual cycle is removed by subtracting the daily climatology over the record (1979-2019). A Fast Fourier Transform high-pass filter is then applied to the z500 anomalies to remove seasonal oscillations (frequencies smaller than $\frac{1}{120days}$) to ensure the network focuses on subseasonal anomalies. The median of the z500 anomalies for the training data (see 2.2.1) is subtracted to obtain an equal number of positive and negative values. These anomalies are then converted into 0s and 1s depending on the sign (negative or positive, respectively). To filter the testing data, z500 anomalies from 2017-2019 are appended to the unfiltered z500 anomalies from 1979-2016 and another FFT high pass filter is applied to all years. The now filtered 2017-2019 data are then subset and used as testing data. The median of the z500 anomalies for the training data (see 2.2.1) is then subtracted and the anomalies are converted into 0s and 1s. For OLR, the annual cycle is removed by subtracting the first 3 harmonics of the daily climatology from the raw field. The first 3 harmonics are used instead of the daily mean because OLR is a noisier field than z500.

110

2.2 Methods

111

2.2.1 Artificial Neural Network Architecture

112

113

114

115

116

117

118

119

120

121

122

123

124

125

126

127

128

129

A two layer ANN (Figure 1) is tasked to ingest tropical OLR and predict the *sign* of the z500 anomaly over the North Atlantic (40°N, 325°E; red dot in Figure 1) 22 days later. The North Atlantic is chosen for this analysis since the MJO is known to force circulation anomalies over this region on subseasonal timescales and thus allows us to explore the utility of an ANN in the context of a well known problem (e.g. Cassou, 2008; Roundy et al., 2010; Henderson et al., 2016). In addition, we find that this grid point is representative of a larger area within the North Atlantic (see supplemental Figure S1).

Each input sample to the ANN consists of vectorized daily anomalous OLR from 30°N to 20°S and 45 to 210°E, where the number of input nodes is equal to the number of OLR grid points ($N = 1407$). The ANN then outputs two values that describe the categorical prediction, positive or negative sign of z500, given the initial OLR input image. The softmax activation function is applied to this final layer and transforms the two output values such that they sum to 1. The output then represents an estimation of the likelihood that an input belongs to a particular category. We refer to this estimation of likelihood as “model confidence”. A more confident prediction will, therefore, have a predicted category value closer to 1. We define forecasts of opportunities as the top 10% most confident predictions by the network, although we explore alternative percentages as well.

130

131

132

133

134

135

136

137

138

139

140

141

142

The ANN architecture consists of two hidden layers of 128 and 8 nodes, respectively, and both use the rectified linear activation function. The final layer includes 2 nodes and uses the softmax activation function. Categorical cross entropy is used for the loss function. This architecture is chosen because it was found to consistently lead to reasonably high accuracies across many combinations of training/validation sets, but our ANN approach should be equally applicable to both shallow and deep networks. The batch size is set to 256 samples (i.e. OLR vectorized images) and the ANN is trained for 50 epochs unless the validation loss increases for two epochs in a row. If this occurs, the ANN stops training early and restores the model’s best weights to reduce overfitting. It is found that 50 epochs is sufficient for training as the ANN rarely completes all 50 epochs. A more detailed explanation of ANNs is provided in the supplemental material for reference along with a comparison of this ANN approach to multinomial logistic regression.

143

144

145

146

147

148

149

150

151

152

The data used to train and test the ANN is composed of three groups: training, validation, and testing. Training and validation data are used during training, where training data is used to update the weights and biases of the ANN and the validation data is used to evaluate the model. The testing data is data that has never been “seen” by the ANN to evaluate the ability of the ANN to generalize to new data. To create the testing data, we assume that the years 2017-2019 have not yet occurred when training the model. In this way, these years act as true testing data for the ANN. While the specific accuracies likely would change with different testing data, the main point of this paper is to introduce a method to identify forecasts of opportunity and then to further identify the associated relevant regions for the enhanced prediction skill, not to provide the most accurate model for this scenario.

153

154

155

156

157

158

159

160

161

162

For this analysis, the ANN validation data is from November 2007 through February 2011 ($N = 481$) and the testing data is from November 2017 through February 2019 ($N = 240$). The remaining extended boreal winter (NDJF) data are used for training (November 1979 - February 2007 and November 2011 - February 2016; $N = 4450$; see supplemental Figure S2). All data is standardized for each grid point by the years used for training and validation. To choose a model for the following analysis, ANN training is repeated for a variety of validation years. Different consecutive four-year chunks are removed from the training data and set aside to use as validation. For each of the nine four-year chunks, the ANN was trained 20 times with random initialized weights. We find that our conclusions are robust to our choice in training period and do not change with variations

163 in random initialization weights. We present one model with reasonably high accuracy
 164 here and using the training, validation, and testing groups outlined above.

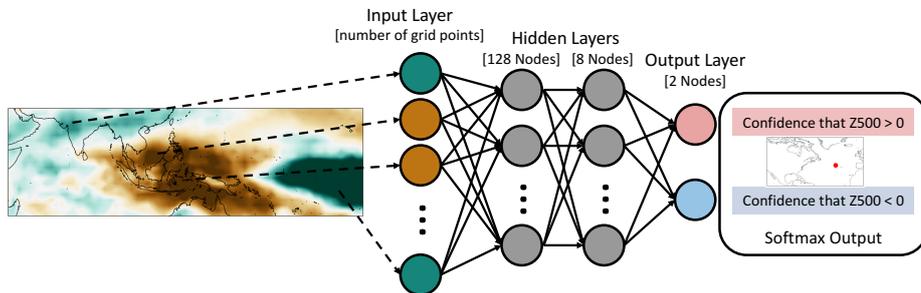


Figure 1. Artificial neural network architecture for prediction of the sign of z500 anomalies over the North Atlantic 22 days following tropical OLR anomalies. The neural network consists of two hidden layers of 128 and 8 nodes, respectively, and an output layer of two nodes (one node for each sign). The output layer uses the softmax activation function.

165 **2.2.2 Layer-Wise Relevance Propagation (LRP)**

166 While ANNs are a useful tool for making predictions, in doing so, they are learn-
 167 ing *how* to make accurate predictions. Therefore, understanding the inner workings of
 168 a trained ANN can provide valuable information for improving prediction skill and un-
 169 derstanding, as well as increasing user confidence in the results. Here, we utilize a re-
 170 latively new neural network explainability technique to the geosciences called layer-wise
 171 relevance propagation (LRP; Bach et al. (2015); Montavon et al. (2019)) to extract and
 172 visualize the features the trained ANN employs to make accurate predictions. While Toms
 173 et al. (2020) describes the use of LRP for geoscience applications in detail, we briefly pro-
 174 vide a high-level description here (see supplemental material for a more detailed expla-
 175 nation). After network training is completed, a single sample is passed through the net-
 176 work and a prediction is made (in our case, two output values are predicted). Our im-
 177 plementation of LRP then takes the highest of these values (i.e. the winning category)
 178 and back-propagates this value through the network via a series of predefined rules, ul-
 179 timately distributing it across the input nodes (i.e. input gridpoints). What results is
 180 a heat map of “relevance” across the input space, where input nodes that are more rel-
 181 evant for the network’s specific prediction for that sample are given higher relevance. This
 182 process is then repeated for every sample of interest, resulting in a unique relevance heat
 183 map for each sample. In our study, since the input layer consists of maps of OLR anom-
 184 alies, the LRP heat maps are maps of the relevant tropical OLR patterns for each pre-
 185 diction of the circulation in the North Atlantic (40°N, 325°E). These maps are discussed
 186 in detail in Section 3.2.

187 **3 Results**

188 **3.1 Identifying Forecasts of Opportunity**

189 ANNs with the architecture shown in Figure 1 are trained 100 times with random
 190 initialized weights to predict the sign of the z500 anomalies 22 days following the trop-
 191 ical OLR anomalies. Figure 2a shows the distribution of the testing prediction accuracy
 192 for all 100 models, where dark teal represents the distribution of all predictions and light

193 teal represents the distribution of the 10% most confident predictions. The correspond-
 194 ing colored vertical dashed lines indicate a threshold for what is expected by random chance.
 195 To calculate the random chance accuracy threshold, 100,000 randomly generated groups
 196 (N=240 for all and N=24 for 10% most confident predictions) of zeros and ones are used
 197 to create a distribution of accuracies, and the 90th percentile of this distribution is used
 198 as the random chance threshold. In Figure 2a, the top 10% most confident prediction
 199 accuracies (light teal) are shifted towards higher accuracies compared to the distribu-
 200 tion with all predictions (dark teal). This shift in the distributions demonstrates that
 201 in general, higher model confidence leads to substantially enhanced prediction accuracy.

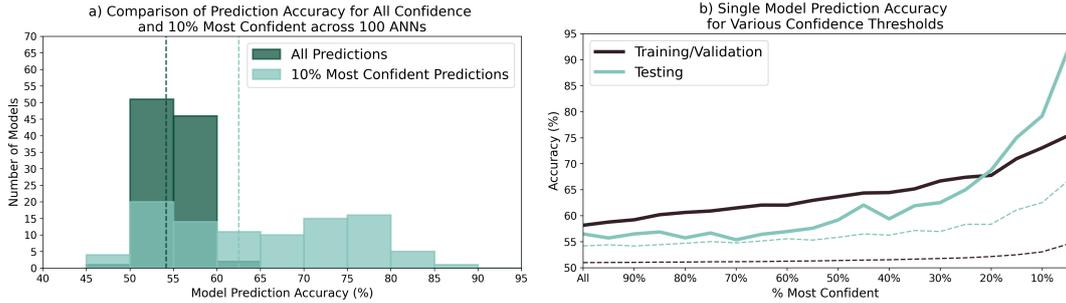


Figure 2. (a) Histograms of testing prediction accuracy for 100 trained ANNs. The dark teal represents the histogram of all prediction accuracies and the light teal represents the histogram for the 10% most confident prediction accuracies. The dark teal and light teal dashed lines in (a) are the maximum accuracies expected by random chance at the 90% confidence level for the corresponding colored histogram (see text for details). (b) Accuracy of one particular model as a function of the percent most confident predictions for training and validation (black) and testing (light teal) data. The dashed lines indicate the maximum accuracies expected by random chance at the 90% confidence level for the corresponding colored lines (see text for details).

202 We chose one model from Figure 2a to further understand how accuracy varies when
 203 a different percent model confidence is used (Figure 2b). The solid lines represent the
 204 accuracy across various model confidence values for training and validation (black) and
 205 testing (light teal) data sets. Figure 2b shows that the testing accuracy (light teal line)
 206 barely outperforms the random chance 90% confidence bound (light teal dashed line) for
 207 all predictions (“all”) while the skill is substantially larger than random chance for the
 208 top 10% of predictions. Accuracy increasing with increasing model confidence is also ap-
 209 parent in the training and validation data. Together, Figure 2a and b illustrate that model
 210 confidence and prediction accuracy generally increase together and therefore, can be used
 211 to identify forecasts of opportunities, or periods of enhanced prediction skill. From this
 212 analysis, the 10% most confident predictions are chosen to define forecasts of opportu-
 213 nity since this threshold has one of the largest accuracy differences from random chance
 214 while still retaining 10% of the samples.

215 When evaluating the network with the training and validation data, the prediction
 216 accuracy for all predictions is 58% and for the top 10% most confident predictions is 73%.
 217 For the testing data, the prediction accuracy for all predictions is 56% and for the top
 218 10% most confident predictions is 79%. The ANN predictions as a function of time are
 219 detailed in Figure S2, and additional skill metrics are provided in Figure S3 and Table
 220 S1.

221

3.2 Tropical Sources of Predictability

222

223

224

225

226

227

228

229

230

231

232

233

We have shown that ANNs can identify forecasts of opportunity using model confidence; however, understanding where this enhanced skill originates is critical for improving physical understanding as well as gaining trust in the network's predictions. To do so, layer-wise relevance propagation is used to identify the OLR patterns that lead the ANN to make correct predictions (see Section 2.2.2). The correct 10% most confident predictions from the training, validation and testing data sets are combined for this LRP analysis. All three sets of data are used instead of only testing data because all data sets have similar accuracies and LRP values (not shown). Thus, including all the data increases the sample sizes for the analysis. The shading in Figure 3c-h shows the regions the network found relevant, on average, to make confident and correct positive (Figure 3c,e,g) and negative (Figure 3d,f,h) z500 predictions. The contours correspond to the average OLR anomalies for these confident and correct predictions.

234

235

236

237

238

239

240

241

242

The average LRP heat map for the correct forecasts of opportunity of positive sign predictions (Figure 3c) indicates four hot spots, one over the southern Indian Ocean into the southern Maritime Continent (20-0°S, 70-130°E), one over the western Pacific (20-0°S, 155°E-170°E), another northwest of Hawaii (25°N, 170°W), and the fourth over Saudi Arabia (30°N, 40-60°E). The average LRP heat map for the correct forecasts of opportunity of negative sign predictions (Figure 3d) indicates four hot spots, one over the Maritime Continent (20-0°S, 110-150°E), one in the western and central Pacific Ocean (20-0°S, 155°E-170°W), another to the west of Hawaii (20°N, 170°W), and the fourth over Saudi Arabia (30°N, 40-60°E).

243

244

245

246

247

248

249

250

251

252

253

254

255

256

For both sign predictions, the hot spots over the Maritime Continent and the western Pacific have opposing signed OLR anomalies (contours) that straddle 150°E. These dipoles of convection over the Indian Ocean into the Maritime Continent and over the western Pacific have similar structures to phase 4-5 and phase 1,7-8 of the MJO (Wheeler & Hendon, 2004). This structure of OLR is consistent with previous research of MJO teleconnections over the North Atlantic for average lead times of 10-14 and 15-19 days (e.g. Cassou, 2008; Henderson et al., 2016; Henderson & Maloney, 2018; Tseng et al., 2018). In addition, this dipole structure is known to lead to higher pattern consistency of teleconnections in the midlatitudes (Tseng et al., 2019), which has been shown to lead to enhanced prediction skill (Tseng et al., 2018). Rossby waves initiated by the MJO tend to be quasi-stationary, which suggests that these OLR anomalies may also correspond to 22 day leads as well. This Maritime Continent and western Pacific Ocean dipole highlighted in part by LRP is therefore consistent with previous research and demonstrates that the ANN has learned physically relevant structures.

257

258

259

260

261

262

263

264

265

266

To test the robustness of these average LRP results for this particular ANN, we calculated the frequency of occurrence of average relevance hotspots greater than 0.5 for models with testing accuracies greater than 70% (Figure 3a,b, $n = 42$ models). We find that all of the hotspots (i.e. the MJO-like structure, the hot spot over Saudi Arabia and the hot spot west of Hawaii) are robust features for enhanced subseasonal prediction throughout these 42 models. In the next section, we hypothesize that the hot spot over Saudi Arabia is associated with the two-way relationship between the North Atlantic Oscillation (NAO) and the MJO (Lin et al., 2009). On the other hand, the hot spot west of Hawaii in both sign predictions is discussed as a possible new region relevant for enhanced subseasonal prediction.

267

3.2.1 *K-means Clustering of LRP Maps*

268

269

270

271

272

273

To further distinguish the relevant regions for the ANN's predictions, k-means clustering (Hartigan and Wong (1979), see supplemental material for more information) is applied to the LRP maps (Figure 3e-h). This analysis reveals that the composite LRP maps for each sign (Figure 3c,d) actually consist of multiple distinct patterns used by the ANN. For positive sign predictions (Figure 3e,g), both clusters have a hot spot located between the central Indian ocean and the maritime continent, which are associ-

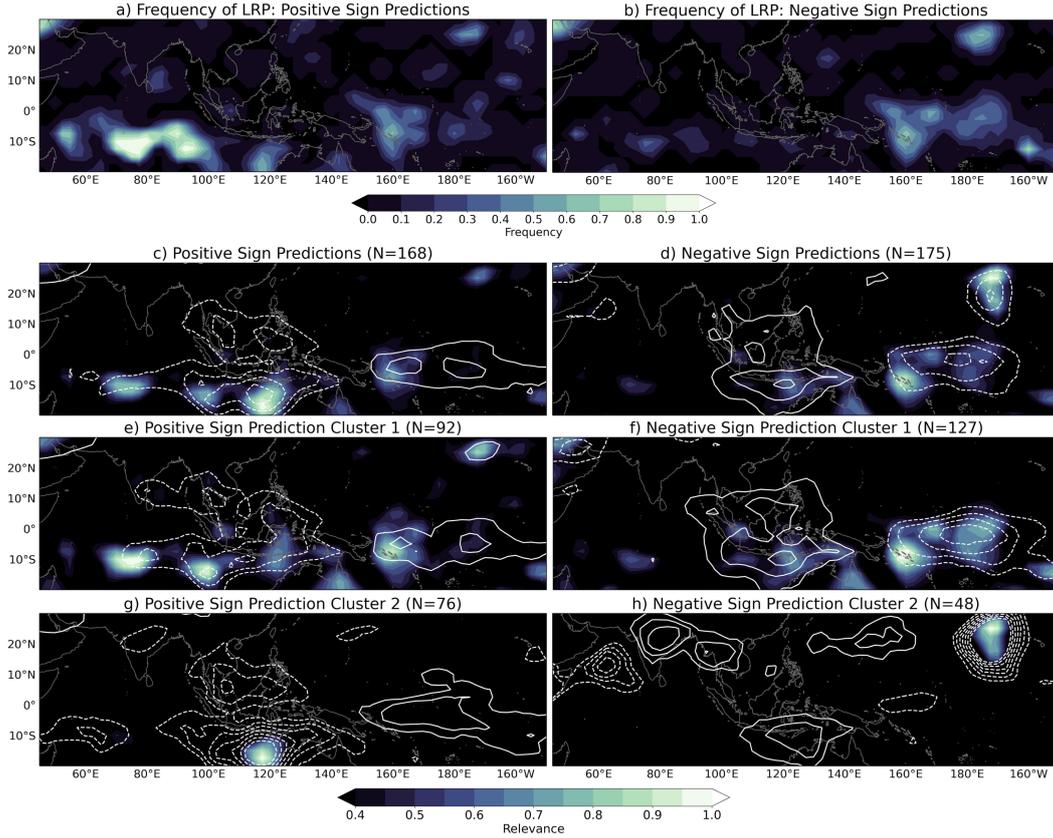


Figure 3. (a,b) LRP frequency of occurrence maps for average relevance values greater than 0.5. Both (a) and (b) consist of models from every 4-year validation chunk. Of these models, only average LRP maps of confident and correct predictions (training, validation, and testing) from models with testing accuracies greater than 70% are included. Maps (c-h) are the LRP maps associated with the ANN from Figure 2b where the shading denotes smoothed composites of LRP fields for all correct forecasts of opportunity for (c) positive sign and (d) negative sign predictions across training, validation and testing periods. The associated two k-means clusters of LRP for (e,g) positive sign predictions and (f,h) negative sign predictions are also shown. Contours represent the corresponding smoothed OLR anomalies where solid lines are positive values and dashed lines are negative values. (a) and (b) contours range from $0.4-1.0 \frac{W}{m^2}$ and $-1.0- -0.4 \frac{W}{m^2}$ and (e-h) contours range from $0.4-1.6 \frac{W}{m^2}$ and $-1.6- -0.4 \frac{W}{m^2}$, both with a contour interval of $0.2 \frac{W}{m^2}$.

274 ated with negative OLR anomalies. While not highlighted by LRP in cluster 2 (Figure 3g),
 275 each negative OLR anomaly region is accompanied by a region of positive OLR anom-
 276 alies over the western Pacific. This suggests the model is identifying an MJO-like pattern.
 277 More specifically, the clustering has identified two types of relevance for this MJO-like
 278 pattern. The LRP map for cluster 1 (Figure 3e) highlights both the positive and nega-
 279 tive OLR anomalies. As previously mentioned, these regions lead to more consistent
 280 midlatitude teleconnections (Tseng et al., 2018) and have been shown to be associated
 281 with a positive NAO anomaly (Cassou, 2008), which corresponds to a positive z500 anomaly
 282 at the predicted location. Cluster 1, therefore, supports previously identified tropical OLR
 283 regions and patterns ideal for enhanced prediction skill on subseasonal timescales in the
 284 North Atlantic. On the other hand, the LRP map for cluster 2 (Figure 3g) focuses ex-
 285 clusively on the south-central Maritime Continent, which is associated with enhanced

286 convection from the Indian Ocean to the Maritime Continent. This is more consistent
 287 with recent research that suggests that convection over the Indian Ocean dominates the
 288 formation of a positive NAO anomaly (Shao et al., 2020). This relationship is nicely il-
 289 lustrated in Figure 3a as the Indian Ocean is highlighted by the LRP analysis more of-
 290 ten than the western Pacific.

291 For cluster 1 of the negative sign predictions (Figure 3f), there are two hot spots,
 292 one over the Maritime Continent and the other over the Pacific Ocean. As with the pos-
 293 itive sign predictions, each hot spot is associated with opposing sign OLR anomalies, how-
 294 ever, unlike cluster 1 of the positive sign predictions, the LRP analysis more strongly high-
 295 lights the western Pacific region, and suggests that the network finds the region of en-
 296 hanced convection more relevant. This is similar to cluster 2 of the positive sign predic-
 297 tions and is consistent with Figure 3b which shows that the region over the western Pa-
 298 cific is more often highlighted by the LRP analysis compared to the Maritime Continent.
 299 This suggests that the network often focuses on the region of enhanced convection for
 300 both sign predictions.

301 Unexpectedly, there is also a hot spot located over Saudi Arabia in cluster 1 for
 302 both positive and negative predictions. As seen in Figure 3a and b, this region is frequently
 303 highlighted by LRP in many ANNs. This hot spot appears to only be important when
 304 an MJO-like dipole structure is present. To the authors' knowledge, this region has not
 305 been shown to be important for tropical-extratropical teleconnections to the North At-
 306 lantic. However, previous research has shown that there is a two-way relationship be-
 307 tween the MJO and NAO. Following the NAO, there tends to be a significant modula-
 308 tion of the tropical upper troposphere zonal wind over the Atlantic-Africa region (Lin
 309 et al., 2009). This modulation has been hypothesized to play a role in MJO initializa-
 310 tion (Lin et al., 2009; Lin & Brunet, 2011). Since the NAO can persist over many weeks,
 311 the network may be identifying an influence of the NAO on the MJO and back to the
 312 NAO. We leave a deeper exploration of this possible mechanism to future work.

313 Unlike the other clusters, cluster 2 of the negative sign predictions (Figure 3h) has
 314 only one hot spot west of Hawaii (25°N , 170°W) and no MJO-like OLR anomalies. We
 315 hypothesize that this region is physically important as it is located south of the subtrop-
 316 ical jet exit region and is associated with a large OLR anomaly. Rossby waves can be
 317 generated through advection of vorticity by upper level divergence or convergence as-
 318 sociated with OLR anomalies (Sardeshmukh & Hoskins, 1988). Since this hot spot re-
 319 gion is close to the jet exit region, these waves can more easily propagate into the mid-
 320 latitudes or become trapped within the North Atlantic jet and directed into the North
 321 Atlantic (Hoskins & Karoly, 1981; Hoskins & Ambrizzi, 1993). Based on these known
 322 tropical-extratropical teleconnection dynamics, it is likely that this hot spot west of Hawaii
 323 is a new pattern identified by the ANN. This hot spot is also weakly apparent in clus-
 324 ter 1 of the positive sign predictions (Figure 3e), but is associated with MJO-like OLR
 325 anomalies. Given the lack of MJO-like patterns in cluster 2 of the negative sign predic-
 326 tions for this region, we hypothesize that this hot spot in cluster 1 of the positive sign
 327 prediction may not actually be associated with the MJO, but instead acting as an ad-
 328 ditional source of predictability.

329 4 Conclusions

330 Improving subseasonal prediction accuracy and understanding requires identify-
 331 ing opportunities that can lead to enhanced predictability (e.g. Mariotti et al., 2020).
 332 Here, we show that an artificial neural network can identify forecasts of opportunity for
 333 subseasonal prediction using the network's confidence in its prediction. In addition, we
 334 demonstrate that layer-wise relevance propagation can extract knowledge gained by the
 335 ANN to identify relevant physical tropical features important for the predictions. K-means
 336 clustering of the LRP maps further provides insight into multiple distinct patterns used
 337 by the ANN for enhanced prediction and reveals a possible new forecast of opportunity
 338 for prediction over the North Atlantic.

339 The hot spots identified by the ANN provide a stepping stone to further our un-
 340 derstanding of tropical-extratropical teleconnections. For example, lagged composite anal-
 341 ysis or simplified models can be used to further explore the physical mechanisms behind
 342 enhanced midlatitude predictability associated with these regions. In addition, analy-
 343 sis of the incorrect predictions made by the ANN may also be useful for improving our
 344 understanding of ideal tropical patterns for enhanced subseasonal prediction. Finally,
 345 while our application is focused on subseasonal prediction, the approach outlined here
 346 should be applicable to predictions across timescales. Ultimately, this paper demonstrates
 347 that ANNs are not only a useful tool for prediction, but can also be used to gain phys-
 348 ical insight into predictability and subsequently, improve prediction skill.

349 Acknowledgments

350 This research is partially funded by the National Science Foundation Graduate Research
 351 Fellowship under the grant number 006784 supporting Kirsten J. Mayer and partially
 352 funded by the National Science Foundation Harnessing the Data Revolution through sup-
 353 porting Elizabeth A. Barnes with grant 1934668.

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

355 Data availability: ERA-I reanalysis data are provided by the European Centre for Medium-
 356 Range Forecasts ([https://www.ecmwf.int/en/forecasts/datasets/reanalysis-datasets/era-](https://www.ecmwf.int/en/forecasts/datasets/reanalysis-datasets/era-interim)
 357 [interim](https://www.ecmwf.int/en/forecasts/datasets/reanalysis-datasets/era-interim); Dee et al., 2011). The interpolated OLR data is provided by the NOAA/OAR/ESRL
 358 PSL, Boulder, CO, USA (https://psl.noaa.gov/data/gridded/data.interp_OLR.html;
 359 Liebmann and Smith, 1996).

360 References

- 361 Abbot, J., & Marohasy, J. (2014, March). Input selection and optimisation for
 362 monthly rainfall forecasting in queensland, australia, using artificial neural
 363 networks. *Atmos. Res.*, *138*, 166–178.
- 364 Albers, J. R., & Newman, M. (2019, November). A priori identification of skillful ex-
 365 tratropical subseasonal forecasts. *Geophys. Res. Lett.*, *46*(21), 12527–12536.
- 366 Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.-R., & Samek, W.
 367 (2015, July). On Pixel-Wise explanations for Non-Linear classifier decisions by
 368 Layer-Wise relevance propagation. *PLoS One*, *10*(7), e0130140.
- 369 Cassou, C. (2008, September). Intraseasonal interaction between the Madden-Julian
 370 oscillation and the north atlantic oscillation. *Nature*, *455*(7212), 523–527.
- 371 Chen, T., & Chen, H. (1995). Universal approximation to nonlinear operators
 372 by neural networks with arbitrary activation functions and its application to
 373 dynamical systems. *IEEE Trans. Neural Netw.*, *6*(4), 911–917.
- 374 Dee, D. P., Uppala, S. M., Simmons, A. J., Berrisford, P., Poli, P., Kobayashi, S., ...
 375 Others (2011). The ERA-Interim reanalysis: Configuration and performance of
 376 the data assimilation system. *Quart. J. Roy. Meteor. Soc.*, *137*(656), 553–597.
- 377 Goodfellow, I., Bengio, Y., Courville, A., & Bengio, Y. (2016). *Deep learning*
 378 (Vol. 1). MIT press Cambridge.
- 379 Ham, Y.-G., Kim, J.-H., & Luo, J.-J. (2019, September). Deep learning for multi-
 380 year ENSO forecasts. *Nature*, *573*(7775), 568–572.
- 381 Hartigan, J. A., & Wong, M. A. (1979). Algorithm AS 136: A K-Means clustering
 382 algorithm. *J. R. Stat. Soc. Ser. C Appl. Stat.*, *28*(1), 100–108.
- 383 Henderson, S. A., & Maloney, E. D. (2018, July). The impact of the Madden–Julian
 384 oscillation on High-Latitude winter blocking during el niño–southern oscillation
 385 events. *J. Clim.*, *31*(13), 5293–5318.
- 386 Henderson, S. A., Maloney, E. D., & Barnes, E. A. (2016, June). The influence
 387 of the Madden–Julian oscillation on northern hemisphere winter blocking. *J.*
 388 *Clim.*, *29*(12), 4597–4616.

- 389 Hoskins, B. J., & Ambrizzi, T. (1993, June). Rossby wave propagation on a realistic
390 longitudinally varying flow. *J. Atmos. Sci.*, *50*(12), 1661–1671.
- 391 Hoskins, B. J., & Karoly, D. J. (1981, June). The steady linear response of a spheri-
392 cal atmosphere to thermal and orographic forcing. *J. Atmos. Sci.*, *38*(6), 1179–
393 1196.
- 394 Kingma, D. P., & Ba, J. (2014, December). Adam: A method for stochastic opti-
395 mization.
- 396 Liebmann, B., & Smith, C. (1996). Description of a complete (interpolated) outgo-
397 ing longwave radiation dataset. *Bull. Am. Meteorol. Soc.*, *77*, 1275–1277.
- 398 Lin, H., & Brunet, G. (2011, January). Impact of the north atlantic oscillation
399 on the forecast skill of the Madden-Julian oscillation: IMPACT OF NAO ON
400 MJO FORECAST. *Geophys. Res. Lett.*, *38*(2).
- 401 Lin, H., Brunet, G., & Derome, J. (2009, January). An observed connection between
402 the north atlantic oscillation and the Madden–Julian oscillation. *J. Clim.*,
403 *22*(2), 364–380.
- 404 Madden, R. A. (1986). Seasonal variations of the 40-50 day oscillation in the tropics.
405 *J. Atmos. Sci.*, *43*(24), 3138–3158.
- 406 Madden, R. A., & Julian, P. R. (1971, July). Detection of a 40–50 day oscillation in
407 the zonal wind in the tropical pacific. *J. Atmos. Sci.*, *28*(5), 702–708.
- 408 Madden, R. A., & Julian, P. R. (1972, September). Description of Global-Scale
409 circulation cells in the tropics with a 40–50 day period. *J. Atmos. Sci.*, *29*(6),
410 1109–1123.
- 411 Mariotti, A., Baggett, C., Barnes, E. A., Becker, E., Butler, A., Collins, D. C., ...
412 Albers, J. (2020, January). Windows of opportunity for skillful forecasts
413 subseasonal to seasonal and beyond. *Bull. Am. Meteorol. Soc.*
- 414 Mariotti, A., Ruti, P. M., & Rixen, M. (2018, March). Progress in subseasonal to
415 seasonal prediction through a joint weather and climate community effort. *npj*
416 *Climate and Atmospheric Science*, *1*(1), 1–4.
- 417 McGovern, A., Lagerquist, R., Gagne, D. J., Eli Jergensen, G., Elmore, K. L.,
418 Homeyer, C. R., & Smith, T. (2019, November). Making the black box more
419 transparent: Understanding the physical implications of machine learning.
420 *Bull. Am. Meteorol. Soc.*, *100*(11), 2175–2199.
- 421 Montavon, G., Binder, A., Lapuschkin, S., Samek, W., & Müller, K.-R. (2019).
422 Layer-Wise relevance propagation: An overview. In W. Samek, G. Montavon,
423 A. Vedaldi, L. K. Hansen, & K.-R. Müller (Eds.), *Explainable AI: Interpret-*
424 *ing, explaining and visualizing deep learning* (pp. 193–209). Cham: Springer
425 International Publishing.
- 426 Mundhenk, B. D., Barnes, E. A., Maloney, E. D., & Baggett, C. F. (2018, Febru-
427 ary). Skillful empirical subseasonal prediction of landfalling atmospheric river
428 activity using the Madden–Julian oscillation and quasi-biennial oscillation. *npj*
429 *Climate and Atmospheric Science*, *1*(1), 20177.
- 430 Nielsen, M. A. (2015). *Neural networks and deep learning* (Vol. 25). Determination
431 press San Francisco, CA.
- 432 Pegion, K., Kirtman, B. P., Becker, E., Collins, D. C., LaJoie, E., Burgman, R.,
433 ... Kim, H. (2019, October). The subseasonal experiment (SubX): A multi-
434 model subseasonal prediction experiment. *Bull. Am. Meteorol. Soc.*, *100*(10),
435 2043–2060.
- 436 Robertson, A. W., Kumar, A., Peña, M., & Vitart, F. (2015, March). Improving and
437 promoting subseasonal to seasonal prediction. *Bull. Am. Meteorol. Soc.*, *96*(3),
438 ES49–ES53.
- 439 Roundy, P. E., MacRitchie, K., Asuma, J., & Melino, T. (2010, August). Modula-
440 tion of the global atmospheric circulation by combined activity in the Madden–
441 Julian oscillation and the el niño–southern oscillation during boreal winter. *J.*
442 *Clim.*, *23*(15), 4045–4059.
- 443 Sardeshmukh, P. D., & Hoskins, B. J. (1988, April). The generation of global ro-

- 444 tational flow by steady idealized tropical divergence. *J. Atmos. Sci.*, 45(7),
445 1228–1251.
- 446 Shao, X., Straus, D. M., Li, S., Swenson, E., Yadav, P., & Song, J. (2020). Forcing
447 of the MJO-related Indian Ocean heating on the intraseasonal lagged NAO.
- 448 Toms, B. A., Barnes, E. A., & Ebert-Uphoff, I. (2020, September). Physically in-
449 terpretable neural networks for the geosciences: Applications to earth system
450 variability. *J. Adv. Model. Earth Syst.*, 12(9).
- 451 Tseng, K.-C., Barnes, E. A., & Maloney, E. D. (2018, January). Prediction of the
452 midlatitude response to strong Madden-Julian oscillation events on S2S time
453 scales: PREDICTION OF Z500 AT S2S TIME SCALES. *Geophys. Res. Lett.*,
454 45(1), 463–470.
- 455 Tseng, K.-C., Maloney, E., & Barnes, E. (2019, January). The consistency of MJO
456 teleconnection patterns: An explanation using linear Rossby wave theory. *J.*
457 *Clim.*, 32(2), 531–548.
- 458 Vitart, F., Ardilouze, C., Bonet, A., Brookshaw, A., Chen, M., Codorean, C., ...
459 Zhang, L. (2017, January). The subseasonal to seasonal (S2S) prediction
460 project database. *Bull. Am. Meteorol. Soc.*, 98(1), 163–173.
- 461 Vitart, F., Robertson, A. W., & Anderson, D. L. T. (2012, January). Subseasonal
462 to seasonal prediction project: Bridging the gap between weather and climate.
463 *WMO Bull.*, 61(61).
- 464 Wheeler, M. C., & Hendon, H. H. (2004, August). An All-Season Real-Time mul-
465 tivariate MJO index: Development of an index for monitoring and prediction.
466 *Mon. Weather Rev.*, 132(8), 1917–1932.
- 467 White, C. J., Carlsen, H., Robertson, A. W., Klein, R. J. T., Lazo, J. K., Kumar,
468 A., ... Zebiak, S. E. (2017, July). Potential applications of subseasonal-to-
469 seasonal (S2S) predictions. *Met. Apps*, 24(3), 315–325.