

1 **Deep learning classification of potentially severe**
2 **convective storms in a changing climate**

3 **Maria J. Molina¹, David John Gagne¹, Andreas F. Prein¹**

4 ¹National Center for Atmospheric Research, Boulder, Colorado, USA

5 **Key Points:**

- 6 • A convolutional neural network can robustly classify convection in current and fu-
7 ture climates.
8 • Skillful classifications are based on learned thermodynamic and kinematic char-
9 acteristics of thunderstorms.
10 • Labels derived using heuristics can suffice for deep learning, bypassing the creation
11 of human labeled data.

Corresponding author: Maria J. Molina, molina@ucar.edu

Abstract

A convolutional neural network (CNN) was found to skillfully classify potentially severe convection of a future climate based on learned thermodynamic and kinematic thunderstorm features. The CNN was trained to classify strongly rotating storms from a current climate, then evaluated against storms from a future climate (end of 21st century), and found to perform with skill and comparatively in both climates. Strongly rotating storms were of interest because they are more likely to be supercells, a thunderstorm type that has a greater likelihood of producing tornadoes and large hail, which cause billions of losses and dozens of fatalities every year. Despite training with labels derived from a threshold value of a severe thunderstorm diagnostic (updraft helicity), the CNN learned physical characteristics of organized convection and environments that are not captured by the diagnostic heuristic. Interpretability techniques revealed that strongly rotating storms are associated with rotation signatures and thunderstorm updrafts penetrating comparatively drier vertical mid-levels. Results show that simple heuristics can yield skillful results with CNNs and can be used to generate labeled data for supervised learning frameworks. Most importantly, results from this study show that deep learning is capable of generalizing to future climate extremes and can exhibit out-of-sample robustness with proper hyperparameter tuning. As the climate continues to change, and machine learning techniques continue to proliferate in the physical sciences, it is important to ensure that techniques perform skillfully with unseen outliers and climate signals. This study offers evidence that this objective is possible and based on physical signals.

Plain Language Summary

As temperatures and water vapor continue increasing due to climate change, models that were trained using past data may no longer perform with skill. Here we explored whether the performance of a machine learning model was sensitive to a changing climate. The purpose of the machine learning model was to classify thunderstorms into two groups: potentially severe and potentially non-severe thunderstorms. Potentially severe thunderstorms have a greater likelihood of producing tornadoes and large hail, which are a threat to society. Results show that the model was able to classify thunderstorms with skill in both the current and future climate partly due to the architecture of the model. We also explored the reasons behind the model's skill and found that the model learned thunderstorm features and meteorological information. These results provide us with added confidence that machine learning models can learn physical relationships from weather and climate data and remain skillful in a future climate.

1 Introduction

The recent success of convolutional neural networks (CNNs; Fukushima & Miyake, 1982) in Earth science applications is largely due to their ability to capture nonlinear and spatially invariant details among input variables. This class of deep learning models (LeCun et al., 2015) has proven skillful in various atmospheric science tasks, including detection of weather and climate features (Y. Liu et al., 2016; Lagerquist et al., 2019; Biard & Kunkel, 2019; Toms et al., 2019), emulation of complex model processes (Rasp et al., 2018), and prediction of extreme weather and climate phenomena (Gagne II et al., 2019; Zhou et al., 2019; Ham et al., 2019; Sobash et al., 2020). This study focuses on convection over the central and eastern contiguous United States (CONUS), which at extremes can produce severe hazards (e.g., hail and tornadoes) that pose societal danger. CNNs have already proven skillful for classification and prediction of convective storms in the present climate (Gagne II et al., 2019; Jergensen et al., 2020; Lagerquist et al., 2020). However, as the climate continues to warm, some future thunderstorms may be outliers in the baseline climate (Trapp & Hoogewind, 2016), and these extreme events may be more difficult for CNNs to identify. This article explores the ability of CNNs to

62 classify convection of a future climate, along with the physical reasons for the resultant
63 performance.

64 Climate change is altering the large-scale atmospheric landscape over North Amer-
65 ica, resulting in changes to the frequency and intensity of organized convection (K. L. Ras-
66 mussen et al., 2017; Prein et al., 2017). Future changes to thermodynamic and kinematic
67 fields can impact climatological distributions of convection morphology and associated
68 severe hazards (e.g., tornadoes and large hail; Trapp et al., 2007, 2009; Diffenbaugh et
69 al., 2013). Studies have shown a climate change imprint on various aspects of severe thun-
70 derstorms and associated environments (Allen, 2018), including increases in thermody-
71 namic buoyancy and storm frequency (Brooks, 2013; Hoogewind et al., 2017), increases
72 in convective inhibition (Taszarek et al., 2020), more societal exposure (Ashley & Strader,
73 2016), and an eastward geographic shift of environments over the U.S. favorable for se-
74 vere hazards (Gensini & Brooks, 2018). However, discerning the interplay between ther-
75 modynamic and kinematic components on future convection has been more challenging
76 (Brooks, 2013), given that subtle changes to either field can alter the potential of a thun-
77 derstorm to produce severe hazards (Doswell et al., 1996). This complex interplay, and
78 varying seasonal and geographical trends, limit the broader conclusions that can be de-
79 rived from climate studies of severe convective storms (e.g., Brooks, 2013).

80 In current forecasting applications, advancements in delineating storms capable of
81 producing specific hazards have included the development of environmental proxies and
82 composite indices that take kinematic and thermodynamic factors into account (E. N. Ras-
83 mussen, 2003; R. L. Thompson et al., 2003, 2007, 2012; Gropp & Davenport, 2018). Up-
84 draft helicity (UH) is an example of a diagnostic parameter, which estimates the mag-
85 nitude of rotation within a storm’s updraft using vertical wind speeds and vorticity (Kain
86 et al., 2008). Strongly-rotating storms with high magnitudes of UH (e.g., $\geq 75 \text{ m}^2 \text{ s}^{-2}$)
87 have a greater likelihood to be of supercell morphology (Clark et al., 2013; Sobash et al.,
88 2016), a type of thunderstorm that observations have shown to be more likely to pro-
89 duce severe hazards (Bunkers et al., 2006; Duda & Gallus Jr, 2010). Scalar thresholds
90 for UH have been used to classify model simulated convection, with storms that exceed
91 the predetermined threshold classified as severe (Sobash et al., 2011; Molina, Allen, &
92 Prein, 2020). These dichotomous assignments derived from UH have been used in kilometer-
93 scale climate simulations to estimate changes to severe hazards in a future climate (Trapp
94 et al., 2011; Gensini & Mote, 2015). However, the use of a heuristic to delineate non-severe
95 and severe convection can result in incorrect categorizations of storms that fall near the
96 predetermined threshold. UH values representative of severe convection also vary sea-
97 sonally and regionally, based on the climatological environments that drive severe con-
98 vection activity (Sobash & Kain, 2017; Molina, Allen, & Prein, 2020). Recently, Sobash
99 et al. (2020) trained a CNN to forecast severe hazard potential using severe thunderstorm
100 parameters, showing that a CNN can learn from diagnostics. The focus herein lies on
101 evaluating a CNN’s ability to classify convection and its out-of-sample robustness to a
102 future climate.

103 CNNs are a class of deep learning models canonically used for computer vision tasks
104 because of the capability of processing multiple layers of information to detect nonlin-
105 earities and spatial invariances of features (LeCun et al., 1998; Krizhevsky et al., 2012).
106 Various techniques have been developed to prevent deep learning models from overfit-
107 ting and to improve training stability, such as dropout and batch normalization (Srivastava
108 et al., 2014; Ioffe & Szegedy, 2015), which help CNNs generalize relationships among in-
109 put features and increase prediction accuracy. However, explaining the reasons for model
110 skill has been challenging, due to the complex architecture of CNNs that include many
111 trained weights and biases within hidden layers and feature maps. Various CNN tech-
112 niques have been recently developed to increase interpretability of machine learning (Barnes
113 et al., 2019; McGovern et al., 2019). These techniques include saliency maps (Simonyan
114 et al., 2013) and permutation feature importance (Breiman, 2001; Lakshmanan et al.,

2015), which have been shown to help explain skillful CNN predictions of convective hazards (Gagne II et al., 2019). Identifying reliable reasons for model performance can increase the trust of atmospheric scientists in machine learning and foster further discovery of the physical processes driving societally impactful weather and climate extremes.

Using deep learning and interpretation techniques, the following questions will be analyzed in this paper:

1. Are future strongly rotating storms correctly classified by a CNN that was trained under current climate conditions?
2. Which input features and spatial patterns are identified to be most important by the deep CNN for correct classifications?
3. What are the reasons for incorrect classifications?

2 Data and Methods

2.1 Storm Identification in Current and Future Climate Simulations

A set of two convection-permitting model simulations created by the Water System Program of the National Center for Atmospheric Research were used to extract storm objects for this study (C. Liu et al., 2017). The two simulations were created using the Weather Research and Forecasting model (WRF; Skamarock & Klemp, 2008) at 4 km grid spacing over the CONUS. The WRF simulations cover 13 years each and represent a retrospective climate period (October 2000–September 2013) and a future climate period (end of the 21st century). Initial and boundary conditions for both simulations were driven by the 6-hourly and 0.7° ERA-Interim (Dee et al., 2011), which is a global climate reanalysis data set produced by the European Centre for Medium-Range Weather Forecasts. A pseudo-global warming (PGW) perturbation signal (Schär et al., 1996), representative of an end of the 21st century business as usual climate scenario, was added to state variables of the future climate simulation. The PGW signal was derived from a set of 19 Coupled Model Intercomparison Project Phase 5 (CMIP5) models (Taylor et al., 2012) generated with a Representative Concentration Pathway of 8.5 W m^{-2} (RCP8.5) radiative forcing, which is a very high greenhouse gas concentration pathway (Moss et al., 2010). To prevent drifting of the 4 km regional simulation from the reanalysis boundary conditions, large-scale spectral nudging of moderate strength was applied above the planetary boundary layer (von Storch et al., 2000), which provided synoptic-scale fidelity to past weather events yet allowed the mesoscale to evolve with some freedom. These model simulations allow us to isolate thermodynamic signals from kinematic influences on the future climate. Simulation details are available in Table 1 and additional specifications can be found in C. Liu et al. (2017).

The watershed transform (Lakshmanan et al., 2009) was used to identify high-intensity updrafts that constitute thunderstorms from the convection-permitting climate simulations. The watershed transform, as employed herein, identified storms using a simulated radar reflectivity minimum threshold of 40 dBZ, which is a quantity proportional to the number of drops per unit volume and provides an estimate of convective precipitation (Trapp et al., 2011). Grid cells adjacent to the detected local maxima that also exceeded a minimum threshold of 20 dBZ were then treated as a part of the storm patch. This process was repeated iteratively and surrounding grid cells were continually associated with a storm until values were either below a minimum threshold of 20 dBZ or exceeded a predetermined storm patch spatial extent of 128 km (32 grid cells x 4 km grid spacing). Each storm was saved as a patch spanning 128 x 128 km containing the storm and the adjacent environment, which influences storm characteristics (R. L. Thompson et al., 2012). Storms were extracted over land and east of the Rocky Mountains (Fig. 1), where severe thunderstorms have a greater climatological likelihood of occurrence (Brooks et al., 2003). The temporal focus of this study was limited to winter (December, January,

Table 1. WRF simulation parameterization schemes^a and settings, as detailed in C. Liu et al. (2017).

Model specifications	
Domain grid points	1,360 x 1,016 grid points
Domain size (East-West, North-South)	5,440-km, 4,064-km
Vertical levels	51 stretched vertical levels, topped at 50-hPa
Microphysics scheme	Thompson aerosol-aware (G. Thompson & Eidhammer, 2014)
Planetary boundary layer scheme	Yonsei University (Hong et al., 2006)
Shortwave and longwave radiation scheme	RRTMG (Iacono et al., 2008)
Land surface scheme	Improved Noah-MP land-surface model (Niu et al., 2011)

^aNo sub-grid cloud cover, shallow, or deep cumulus parameterizations were employed.

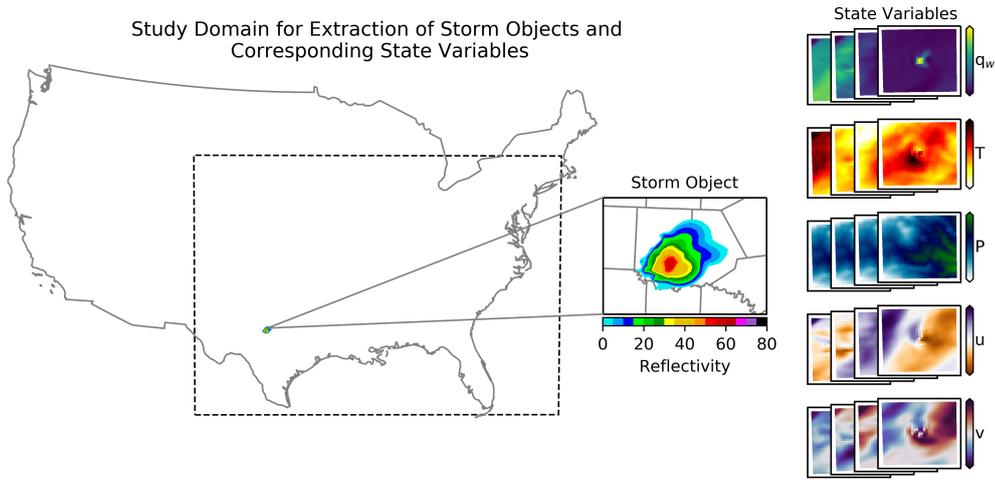


Figure 1. Storm objects for this study were extracted from areas east of the Rocky Mountains (over land) within the dashed-line polygon. An example storm object is shown over the CONUS for scale, with the inset displaying a larger version, and corresponding state variables are shown on the right. State variables listed from top-to-bottom are water vapor mixing ratio (q_w ; g kg^{-1}), temperature (T ; K), pressure (P ; hPa), and zonal (u) and meridional (v) winds (m s^{-1}). The four layers for each state variable indicate the four levels (1, 3, 5, and 7 km above ground) at which variables were derived.

165 and February; DJF) and spring months (March, April, May; MAM). Other seasons were
 166 omitted due to a simulated dry bias during summer months across the central CONUS,
 167 which was partly associated with land-surface feedbacks (Barlage et al., 2018).

168 Similar to Gagne II et al. (2019), meteorological state variables were extracted from
 169 the WRF simulations to train a CNN after creating the storm patches. The variables
 170 are pressure (P ; hPa), temperature (T ; K), water vapor mixing ratio (q_w ; g kg^{-1}), and
 171 zonal (u) and meridional (v) winds (m s^{-1}). Variables were then interpolated onto the
 172 following heights above ground level (AGL): 1, 3, 5, and 7 km. AGL heights were preferred
 173 over constant pressure surfaces because pressure surfaces might be below ground
 174 across portions of the High Plains and AGL heights are more likely to sample similar parts
 175 of a storm updraft. UH ($\text{m}^2 \text{s}^{-2}$) was also extracted and is quantified as

$$\text{UH} = \int_{2\text{km}}^{5\text{km}} w \zeta dz,$$

176 where the integral of the product of vertical velocity (w) and vertical vorticity (ζ) is com-
 177 puted from 2 km to 5 km AGL (Kain et al., 2008). A 1D vector containing labels for CNN
 178 training and testing was created using binary assignment (i.e., one-hot encoding) derived
 179 from UH. A high-magnitude UH threshold (e.g., $75 \text{ m}^2 \text{ s}^{-2}$) was used to delineate con-
 180 vection more likely to be of supercell morphology (Sobash et al., 2011). Values exceed-
 181 ing the UH threshold were assigned a label of 1, whereas values below the threshold were
 182 assigned a label of 0. Storm objects were then split into two subsets prior to CNN train-
 183 ing: 60% for training and 40% for testing. Since the meteorological variables contain dif-
 184 ferent dynamic ranges, the training data was standardized by subtracting the training
 185 set variable’s mean and then dividing by its standard deviation.

186 2.2 CNN Architecture and Interpretation

187 The deep learning model used in this study was a CNN (LeCun et al., 1990) that
 188 consisted of three convolutional layers (similar to Gagne II et al., 2019), which allowed
 189 the model to learn features of various spatial scales to then perform a classification task
 190 (Fig. 2). At each convolutional layer, a 2D filter window was slid across each input map,
 191 which then outputted a “feature map”. The input maps received by the first layer were
 192 the storm patches that contained standardized state variables and the maps received by
 193 deeper layers were the previous layer feature maps (Fig. 2). Each filter pass contained
 194 a set of weights that were randomly initialized before the first training pass. The filter
 195 window was slid across the input map one grid cell at a time in this study (stride length
 196 of 1), with this computation performed iteratively until spanning the full map space (more
 197 details available in Goodfellow et al., 2016, and others). The dimensions of the filter win-
 198 dow were 5×5 grid cells (20×20 km), with zero padding also applied to the edges of
 199 each map in order to obtain an output map of the same size after the convolution op-
 200 eration. A nonlinear activation function was then applied to each feature map to enable
 201 the conditional passage of relevant signals for a particular example. The rectified linear
 202 unit (ReLU; $\max(0, x)$) activation function was used for each feature map, which pre-
 203 served the magnitude of positive signals and negated negative signals during training,
 204 when signals were propagated backward through the network (LeCun et al., 2015).

205 The spatial dimensions of feature maps were reduced after each convolutional layer
 206 via maximum (max) pooling, which downsampled each feature map’s spatial extent by
 207 half (Fig. 2). Max pooling was performed by extracting maximum values of the feature
 208 maps within a sliding filter window, which was of 2×2 dimension (grid cell). Max pool-
 209 ing added spatial invariance and allowed the model to learn higher level features in deeper
 210 layers. After the three convolution and dimensionality reduction operations, the resul-
 211 tant data was flattened into a 1D vector and passed through a densely connected layer
 212 (Fig. 2), with a ReLU activation function applied to the output. The 1D vector was then
 213 passed through a final dense layer, with a sigmoid activation function applied to produce
 214 the model’s classification result as a value between 0 and 1, which was interpreted as the
 215 probability that the input state variables contained a strongly rotating thunderstorm.
 216 The weights of the deep CNN were trained to minimize mean squared error (MSE) us-
 217 ing the Adam optimization algorithm (Kingma & Ba, 2014) via backpropagation. Dur-
 218 ing training, a set of 128 examples (i.e., batch) were randomly pulled from the training
 219 data population, and passed forward through the deep CNN to generate storm classi-
 220 fication probabilities. During backpropagation, partial derivatives of the error with re-
 221 spect to the weights for all hidden layers were calculated, and weights were updated to-
 222 wards values that reduced MSE. A learning rate of 0.0001 was used, which controlled
 223 the magnitude of the weight updates.

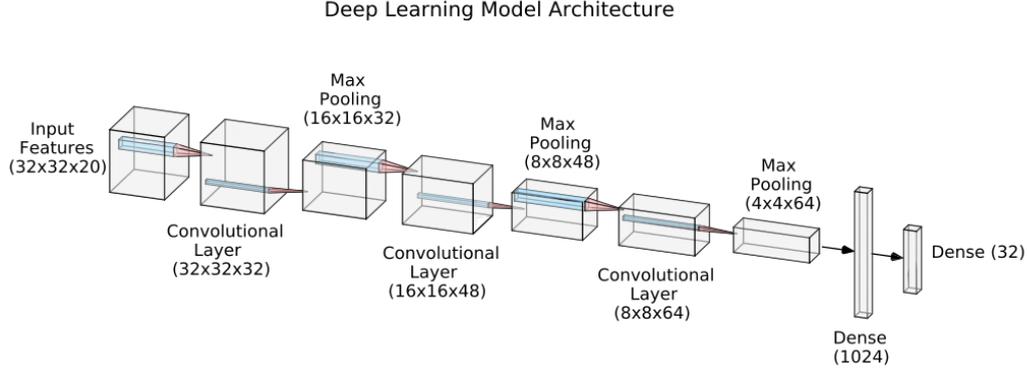


Figure 2. The architecture of the deep convolutional neural network (CNN). The model consists of three 2D convolutional layers and maximum (max) pooling layers. The dimensions of the feature maps (shown in parentheses) decrease in spatial extent and increase in depth as inputs travel deeper into the model, allowing the CNN to learn relationships of coarser spatial resolutions. 2D filter windows are also depicted in blue, of dimension 5×5 for each convolutional layer and 2×2 for each max pooling layer.

224 To prevent overfitting of weights during training, various techniques were employed.
 225 Ridge (L2 norm; 0.001) regularization was added as a penalty term to reduce the mag-
 226 nitude of the weights at each convolutional layer. Batch normalization was applied af-
 227 ter each convolutional layer and the first dense layer (Ioffe & Szegedy, 2015), which in-
 228 volves standardizing layer outputs by subtracting the batch mean and dividing by the
 229 batch standard deviation, in effect reducing covariance shift. This re-centering of fea-
 230 ture maps also increases independence among convolutional layers and reduces the to-
 231 tal training time of the model. 2D spatial dropout (Srivastava et al., 2014) was also em-
 232 ployed after batch normalization, which probabilistically removes feature maps (30% in
 233 this study), reducing the number of features learned and increasing the robustness of in-
 234 dividual maps and weights. A validation data set was used during training, consisting
 235 of 10% of the available training data, which provided insight into the skill of the model
 236 during weight optimization (training). These model settings were selected based on the
 237 lowest resultant test data MSE from a hyperparameter grid search that resulted in over
 238 128 independently trained CNNs. All CNNs were trained with 20 epochs. The classi-
 239 fication output of the lowest MSE CNN was evaluated using probabilistic and nonprob-
 240 abilistic skill metrics that will be further detailed within the results.

241 To explore the relative importance of specific meteorological variables on CNN clas-
 242 sification performance, we used the permutation feature importance (PFI; Breiman, 2001)
 243 analysis. PFI ranks variables based on how much randomizing them impacts error dur-
 244 ing testing, with larger magnitude decreases in skill associated with greater importance.
 245 Higher relative importance suggests that the respective variables have greater relevance
 246 to the classification due to the larger magnitude weights associated with them within the
 247 CNN architecture. 500 permutations were completed for each of the 20 variables to cap-
 248 ture uncertainty associated with shuffling order in PFI. Permuted fields for a set of ex-
 249 amples were also visualized to further explain variable importance results. The chosen
 250 examples consist of cases that were originally classified correctly by the CNN, but switched
 251 to incorrect classifications due to PFI. Certain classified storms that were switched to
 252 incorrect classifications also consistently appeared in larger skill reductions, and these

253 were used to narrow down the subset of storms for visualization. To explain model rea-
 254 soning within a spatial context, image-specific class saliency visualization was used (Simonyan
 255 et al., 2013). Saliency maps were computed using gradients of the CNN output with re-
 256 spect to the input features. Examples from the test data were chosen to explore pixels
 257 that have high saliency with respect to the model probability output.

258 3 Results

259 3.1 Storm Classification in a Future Climate

260 Storms identified within the future climate model simulation contain warmer tem-
 261 peratures and higher moisture content than storms identified within the current climate
 262 model simulation at all vertical levels (1, 3, 5 and 7 km; Table 2), which is consistent with
 263 the applied PGW signal (C. Liu et al., 2017). Table 2 shows that future storms contain
 264 about 1.3 g kg^{-1} more low-level (1 km) water vapor mixing ratio and are about 2.4 K
 265 warmer (1 km) than storms of the current climate (both statistically significant at the
 266 95th percentile confidence level). These results are also consistent with the Clausius-Clapeyron
 267 equation that estimates a 7% increase in saturation vapor pressure per $+1^\circ\text{C}$. Using the
 268 Clausius-Clapeyron equation, water vapor mixing ratio of future storms should be about
 269 8.76 g kg^{-1} at 1 km, which is comparable to the 8.9 g kg^{-1} contained in storms extracted
 270 from the future climate model simulation (Table 2). Added low-level moisture and warmth
 271 provide additional thermodynamic buoyancy and vertical instability that could lead to
 272 more intense convection in the future (K. L. Rasmussen et al., 2017; Prein et al., 2017).
 273 The increased moisture and warmth could also pose the CNN with added difficulty in
 274 performing the storm classification task. Table 2 shows little change in zonal (u) and merid-
 275 ional (v) storm winds between the current and future climate model simulations. Since
 276 the classification task being performed by the CNN is related to winds, the relative con-
 277 sistency in wind magnitude may result in little change in classification skill between the
 278 current and future climate model simulations.

279 Here we evaluate probabilistic forecasts generated by the CNN, which are prob-
 280 abilities that the storm patches contain a strongly rotating or non-strongly rotating storm.
 281 Strongly rotating storms are associated with a higher probability magnitude and non-
 282 strongly rotating storms are associated with a lower probability magnitude. The large
 283 imbalance between the majority and minority classes was important to consider during
 284 evaluation of the CNN classification skill (Table 3). Therefore, the performance diagram
 285 and metrics that are more useful for evaluating correct forecasts of rare events were used
 286 for evaluation (Roebber, 2009). The minority class in this case consists of strongly ro-
 287 tating storms, which are rare events that comprise approximately 3% of all storms in the
 288 convection-permitting model simulations. Performance diagrams summarize the prob-
 289 ability of detection (POD; ratio of hits to the total of hits and false alarms), critical suc-
 290 cess index (CSI; ratio of hits to the total of hits, false alarms, and misses), and bias (ra-
 291 tio of false alarms to misses). Success ratio (SR) is also summarized, which is $1 - \text{false}$
 292 $\text{alarm ratio (FAR; ratio of false alarms to the total of hits and false alarms)}$. The curves
 293 shown on the performance diagrams were created by varying the probability threshold
 294 between 0 and 1 to convert probabilistic forecasts into binary forecasts and show how
 295 skill changes based on the probability threshold used (Fig. 3a,c,e).

296 The performance diagrams (Fig. 3a,c,e) show that despite being trained with storm
 297 patches extracted from the current climate model simulation, CNN skill remains con-
 298 sistent and high (0.71 max CSI) when classifying storms of the future climate model sim-
 299 ulation (Fig. 3c). These results suggest that a CNN is capable of learning spatial rep-
 300 resentations and variable relationships that are transferable to a warmer and more moist
 301 climate. Figure 4a shows that correctly classified storms of the future climate contained
 302 approximately 4 g kg^{-1} more low-level moisture (1 km) than correctly classified storms
 303 of the current climate. The consistency in CNN skill could be partly related to bulk wind

Table 2. Median of storm variables extracted from the current and future climate simulations. Environments surrounding the storms were omitted for these statistics. Future storms with higher low-level moisture content than most cases in the future climate (i.e., outlier cases with $\geq 99^{th}$ percentile of 1 km water vapor mixing ratio in the future climate), are also shown. Statistically significant values of the future climate and future outliers are indicated in **boldface** and computed using confidence intervals of 2.5th and 97.5th percentile of a 1,000-member bootstrap from a total sample of 454,242 storm objects extracted from the current climate simulation.

Current Climate	1-km	3-km	5-km	7-km
Temperature (K)	283.7	272.7	261.0	247.4
v-winds (m s ⁻¹)	5.3	7.8	9.7	11.2
u-winds (m s ⁻¹)	3.1	9.9	14.1	17.2
Water vapor mixing ratio (g kg ⁻¹)	7.6	4.8	2.1	0.7
Pressure (hPa)	868.7	679.9	526.4	402.3
Future Climate	1-km	3-km	5-km	7-km
Temperature (K)	286.2	275.6	264.6	251.9
v-winds (m s ⁻¹)	5.1	7.4	9.5	11.3
u-winds (m s ⁻¹)	2.8	10.0	14.4	17.8
Water vapor mixing ratio (g kg ⁻¹)	8.9	5.7	2.8	1.0
Pressure (hPa)	869.0	681.7	529.5	406.3
Future Outliers	1-km	3-km	5-km	7-km
Temperature (K)	295.7	284.7	272.5	260.5
v-winds (m s ⁻¹)	4.8	4.0	3.8	4.2
u-winds (m s ⁻¹)	2.6	7.0	9.8	12.0
Water vapor mixing ratio (g kg ⁻¹)	17.2	8.0	3.5	1.4
Pressure (hPa)	889.2	704.1	551.5	427.0

Table 3. Table contains various skill metrics used for evaluation of CNN performance during the current and future climate. Also shown are the total number of true positive (i.e., hits), false positive (i.e., false alarms), false negative (i.e., misses), and true negative predictions made by the CNN. Future storms that have higher low-level moisture content than most cases in the future climate (i.e., outlier cases with $\geq 99^{th}$ percentile of 1 km water vapor mixing ratio in the future climate), are also shown. Metrics were computed using a 0.5 forecast probability threshold.

Climate	Current	Future	Outlier
True positives	9,089	9,654	569
False positives	1,633	1,149	114
False negatives	2,250	3,284	66
True negatives	441,270	442,380	3,818
AUC	0.90	0.87	0.93
CSI	0.70	0.69	0.76
Hit Rate	0.80	0.75	0.90
Bias	0.95	0.84	1.08
BSS	0.74	0.72	0.76
Resolution	0.02	0.02	0.09
Uncertainty	0.02	0.03	0.12

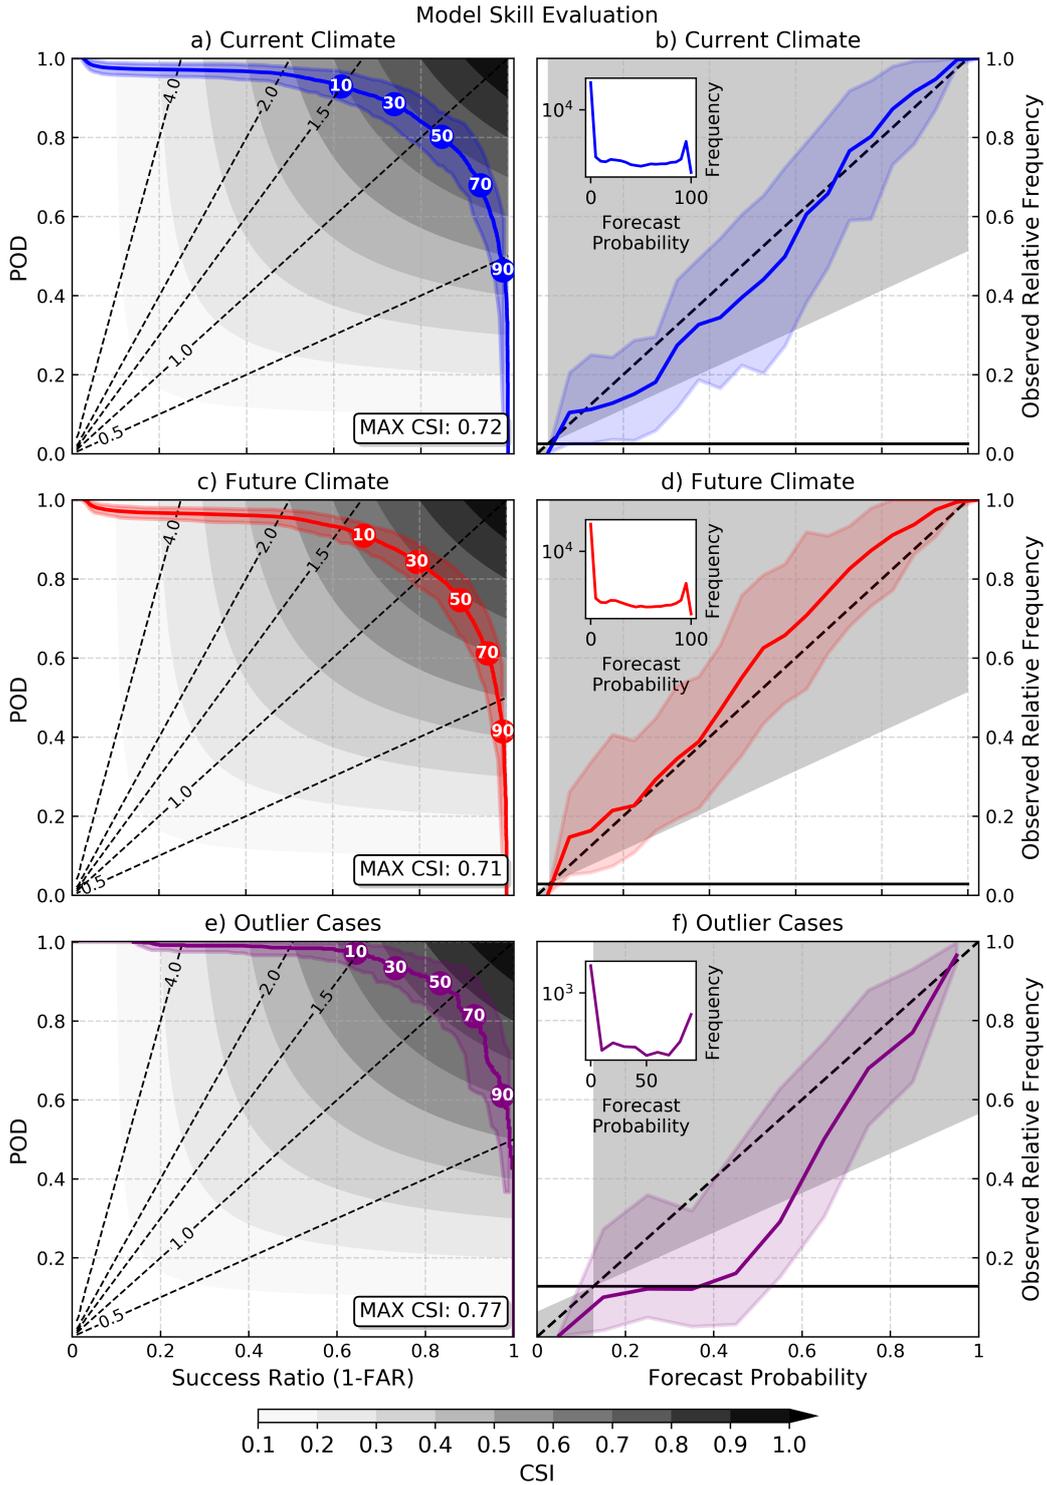


Figure 3. Performance diagrams (a,c,e) show curves that represent CNN skill as a function of the probability of detection (POD) and success ratio (1-FAR [false alarm ratio]) across various probability thresholds. The grayscale filled contours show the critical success index (CSI), the dashed lines display the bias, and circles along the curves display probability thresholds (a,c,e). Attributes diagrams are also displayed, which show forecast probabilities against observed relative frequency, using a forecast probability bin size of 0.05 (b,d) and 0.1 (f). Inset panels in the top left show the frequency of forecast probabilities and the grey-shading shows regions where resolution exceeds reliability (b,d,f). 95th percentile confidence intervals (two-tailed) computed from a 1,000-member bootstrap shown with shading (a-f).

304 shear (1-5 km) distributions that remained relatively stationary between both climate
 305 model simulations (Fig. 4a). Despite the imbalance between the majority and minor-
 306 ity classes, the CNN was able to perform the classification task skillfully, suggesting that
 307 techniques to augment minority classes may not always be necessary (e.g., Chawla et al.,
 308 2002). However, model bias exhibits some sensitivity to the forecast threshold used. Max
 309 CSI and lower bias were achieved when using a probability threshold of approximately
 310 0.35 in the future climate and 0.5 in the current climate. A probability threshold of 0.5
 311 results in an under forecasting bias (<1) of strongly rotating storms of the future (Fig.
 312 3c; Table 3), which shows that the CNN generally has lower confidence in classifying strongly
 313 rotating storms of the warmer and more moist climate.

314 Performance metrics were also computed for future storms that were characterized
 315 by higher low-level moisture content than most cases in the future climate in order to
 316 further quantify the out-of-sample robustness of the CNN (Table 2). These future storms
 317 were classified as “outlier cases” and were selected as storms containing 1 km water vapor
 318 mixing ratio exceeding the 99th-percentile of storms from the future climate (Fig.
 319 3e). The focus of outlier cases lies on 1 km water vapor mixing ratio because increased
 320 low-level moisture and thermodynamic buoyancy can result in more intense vertical winds
 321 related to stronger storm updrafts. Results show that CNN classification skill with out-
 322 lier storms of the future climate remains high (Fig. 3e) and comparable to the current
 323 and future climate subsets (Fig. 3a,c). These results further substantiate that a CNN
 324 can exhibit out-of-sample robustness in climate applications. Results also suggest that
 325 deep learning can sufficiently generalize relationships among input variables and remain
 326 skillful with extreme events. However, some over forecasting of strongly rotating out-
 327 lier storms was identified (bias ≥ 1) with a probability threshold of 0.5 (Fig. 3e; Table
 328 3), which implies overconfidence in classifying storms with extreme low-level moisture.

329 The Brier skill score (BSS) was used as an additional evaluation metric and can
 330 be visualized with the attributes diagram (Fig. 3b,d,f), which show forecast probabil-
 331 ities against observed relative frequencies (Hsu & Murphy, 1986; Wilks, 2011). An at-
 332 tributes diagram provides a measure of forecast reliability (e.g., a storm should be strongly
 333 rotating 60% of the time that a 60% forecast is issued), where the dashed 45 degree line
 334 represents perfect reliability (Wandishin et al., 2005). The solid horizontal line in fig-
 335 ure 3b,d,f shows the climatological probability of strongly rotating storms occurring within
 336 the respective climate sample, which is higher in outlier cases than in the current and
 337 future climates. Since attributes diagrams consider climatological and forecast proba-
 338 bility frequency, they also show how different forecasts are from climatology (i.e., res-
 339 olution). The gray shading in figure 3b,d,f show areas contributing to positive BSS, which
 340 are areas where BSS resolution exceeds reliability (Gagne II et al., 2019). Inset plots (Fig.
 341 3b,d,f) show the frequency of forecast probabilities for each climate subset, which in this
 342 case features a bi-modal distribution, with peaks at low (≤ 0.05) and high (≥ 0.95) fore-
 343 cast probabilities. This bimodal distribution is most pronounced for outlier cases (Fig.
 344 3f). The attributes diagram curves closely parallel the dashed 45 degree dashed diag-
 345 onal line across all forecast probabilities for both the current and future climates (Fig.
 346 3b,d), which conveys high forecast reliability. However, outlier cases have lower reliabil-
 347 ity between the 0.2-0.7 forecast probabilities and higher reliability at low (<0.2) and high
 348 (>0.7) forecast probabilities (Fig. 3f). These results corroborate the performance dia-
 349 gram results, which show that the CNN is overconfident in classifying strongly rotating
 350 outlier storms.

351 3.2 CNN Interpretation

352 Permutation feature importance (PFI) was conducted to determine the relative im-
 353 portance of input variables on CNN prediction skill. The area under the receiver oper-
 354 ating characteristic curve (AUC; Mason, 1982) was used, which is a scalar that repre-
 355 sents model performance encompassing the probability of detection and false detection.

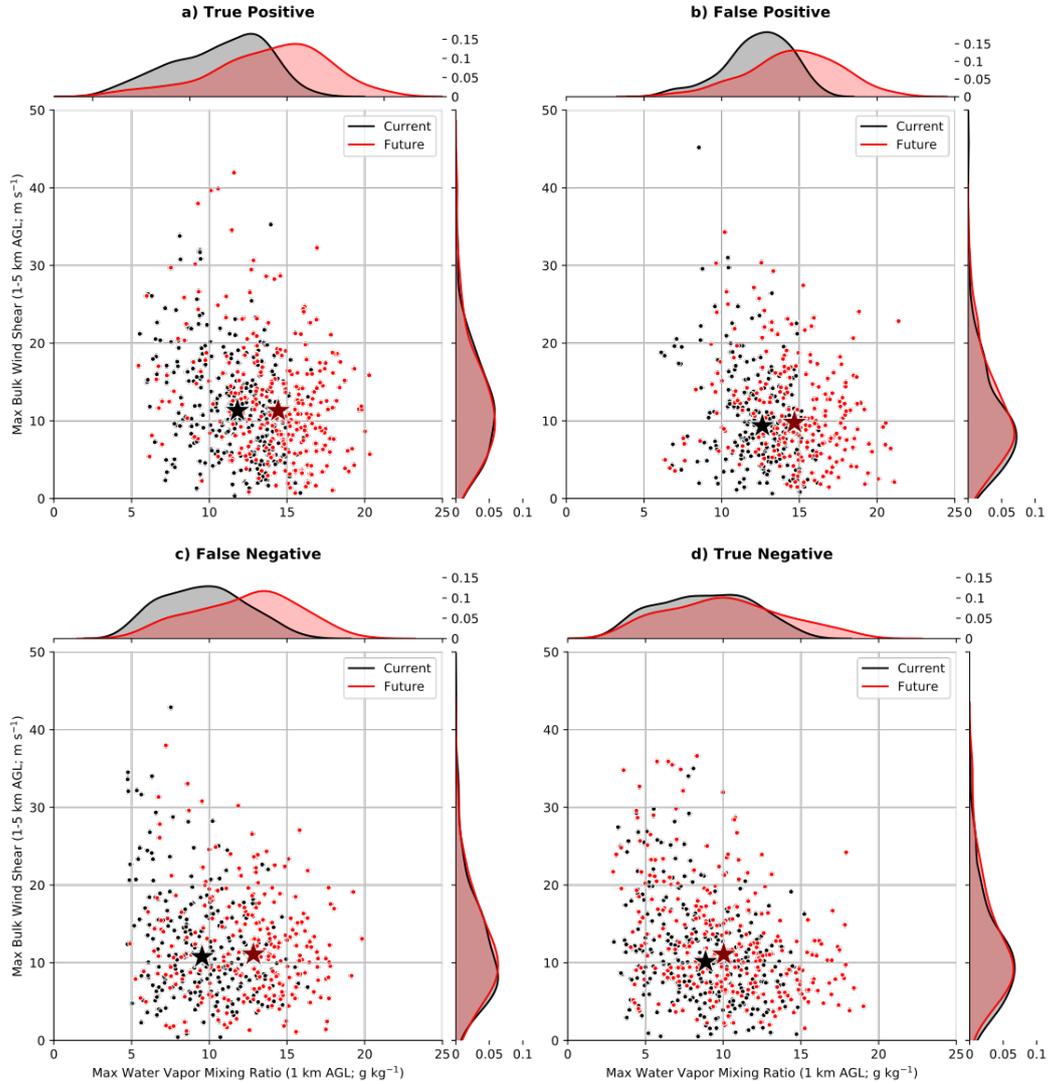


Figure 4. Scatter plots showing water vapor mixing ratio (1 km AGL) against bulk wind shear (1-5 km AGL) for storm objects of the current and future climate evaluated as (a) hits, (b) false alarms, (c) misses, and (d) correct negatives. The dots represent individual storm objects of the current (black) and future (red) climates, while the stars show the mean of the respective climate storm objects. Bivariate density distributions are also shown with marginal plots created using Gaussian kernels. Random subsets of storm objects are shown for easier visualization.

Using AUC, PFI reveals that zonal (u) and meridional (v) winds at 3 km have the highest relative importance for CNN prediction (Fig. 5a,d,g). PFI results are consistent for predictions generated using the current climate, future climate, and future outlier storms (Fig. 5a,d,g), which shows that mid-level kinematic fields play an important role in the proper classification of rotating convective storms. This result is physically reasonable given that UH (computed from 2 km to 5 km AGL) was used to create the storm labels that were subsequently used to train the CNN. The climatological homogeneity between current and future climate mid-level winds (Table 2) also likely contributed to the consistency in variable importance across climate subsets. Zonal and meridional winds at 1 km and 5 km were also identified as important (Fig. 5a,d,g). Several thermodynamic variables also ranked in the top 50th percentile in importance, suggesting that the CNN also relies on characteristics of physical variables that were not included in the UH computation. These relatively higher ranking thermodynamic variables include, temperature at 5 km and water vapor mixing ratio at 1 km and 7 km (Fig. 5a,d,g).

Additional skill metrics were used for PFI in order to explore the sensitivity of the analysis to the respective evaluation method. PFI using CSI, which is a skill evaluation metric that neglects true negative events (as described earlier), further emphasizes the relative importance of mid-level kinematic fields (Fig. 5b,e,h). BSS was also used for PFI (Fig. 5c,f,i) and results generally align with AUC and CSI results in regards to the relative importance of mid-level kinematic fields. Interestingly however, moisture at 5 km ranked most important when evaluating the CNN classification skill for future outlier (Fig. 5h,i) and current climate (Fig. 5c) storms and ranked 4th for future climate storms (Fig. 5f). This result suggests that mid-level moisture is an important variable for correct classification of strongly rotating storms, given the lower ranking found using AUC, which also takes into account correct classification of non-strongly rotating storms. In the case of future outliers, the consistent relative importance of water vapor mixing ratio at 5 km and 7 km across skill metrics, particularly for BSS (Fig. 5g,h,i), is surprising because its future values are substantially above those of the current climate training period (Fig. 6). Figure 6 shows that outlier storms contain 5 km water vapor mixing ratio that is on average 4 g kg^{-1} greater than current climate storms for hits, false alarms, misses, and correct negative cases. This result suggests that the spatial arrangement of meteorological fields likely also plays an important role in CNN prediction skill.

PFI offers insight into the relative importance of variables based on modulations to the CNN prediction skill, but the method does not provide reasons for the rankings. For instance, it is not immediately clear why water vapor mixing ratio at 5 km has greater relative importance than at 1 km. To explore the reasons for PFI rankings, visualizations were created of storms that were initially correctly classified as strongly rotating but incorrectly switched to a non-strongly rotating classification as a result of the permuted variable (Fig. 7). Figure 7c shows an example strongly rotating storm. Its associated water vapor mixing ratio at 5 km (Fig. 7a) was permuted to a field that had a greater overall magnitude of moisture (Fig. 7b) and peak moisture values that were offset from the storm locations (Fig. 7c), which resulted in the incorrect non-strongly rotating classification. Various other storms also had a similar pattern; higher overall moisture content and shifted peak value locations in the permuted field resulted in incorrect non-strongly rotating classifications (not shown). Supercells generally form in environments characterized by moist low-levels and drier mid-to-upper levels, while stratiform precipitation or less organized convection could be characterized by higher and more homogeneous moisture profiles (Bunkers et al., 2006; R. L. Thompson et al., 2012). These examples show that moisture characteristics of vertical atmospheric profiles are likely a learned feature by the CNN. In regards to the high importance of zonal and meridional winds, storms that were incorrectly classified as non-strongly rotating during PFI were due to the uniformity of zonal or meridional winds in the permuted fields (Fig. 7e,h), as opposed to the overall magnitude of the horizontal winds. These results show that

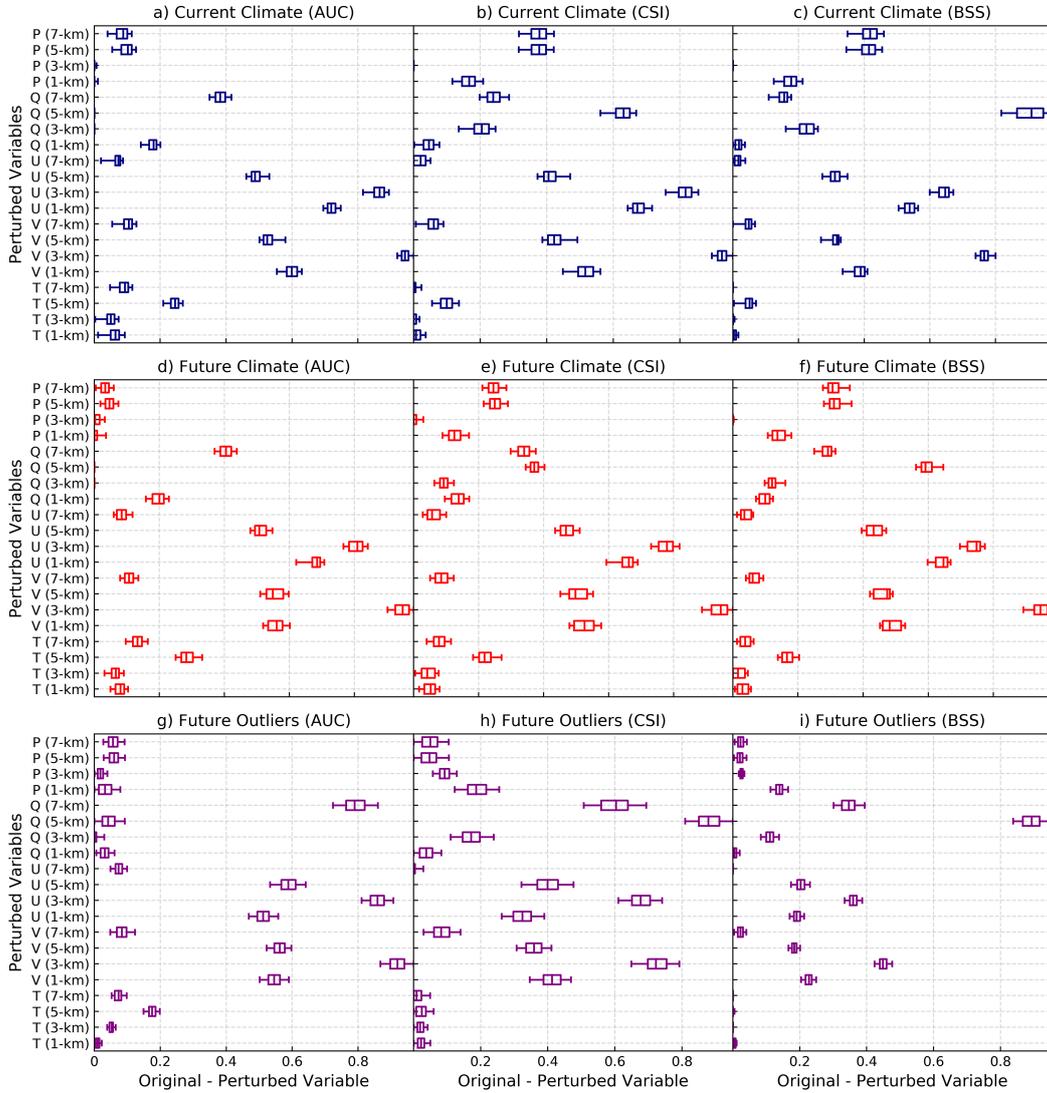


Figure 5. Permutation feature importance (PFI) analysis for the current climate (a,b,c), future climate (d,e,f), and future outlier (g,h,i) storms shown using box and whisker plots. The median of 500 permutations is represented by the vertical line within the box and the whiskers represent all 500 measured changes in skill. PFI was conducted using various skill metrics, including area under the receiver operating characteristic curve (AUC; a,d,g), critical success index (CSI; b,e,h), and Brier skill score (BSS; c,f,i). Changes in skill were normalized by the maximum change in the respective climate subset and skill metric.

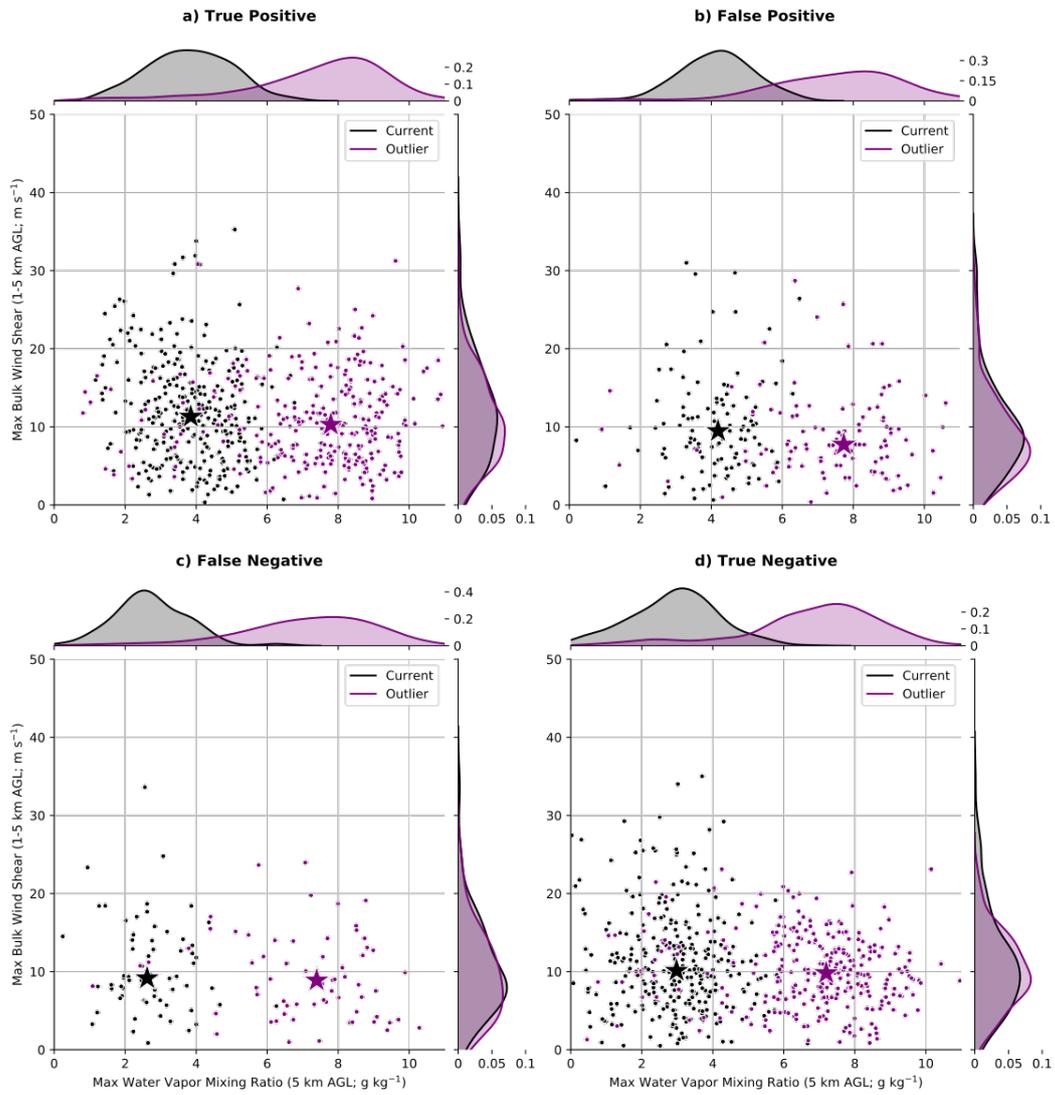


Figure 6. Same as 4, but for outlier storms and water vapor mixing ratio at 5 km AGL.

409 the CNN learned that wind directional shifts over a small region located near the storm
410 core were indicative of strong rotation.

411 Visualizations were also created for storms that were initially correctly classified
412 as non-strongly rotating, but incorrectly switched to a strongly rotating classification dur-
413 ing PFI (Fig. 8). The permuted moisture field was drier and contained large magnitude
414 gradients in space that represented isolated and intense convection (Fig. 8b). Regard-
415 ing kinematic fields, the original zonal (Fig. 8d) and meridional (Fig. 8g) winds lacked
416 rotational characteristics for the respective storms (Fig. 8f,i). However, the permuted
417 fields contained strong rotational features (Fig. 8e,h) which likely resulted in the changed
418 classification.

419 Individual storms from the future climate model simulation were chosen to visu-
420 alize areas of saliency for predictions made by the CNN. Simulated radar reflectivity of
421 the respective examples are shown in figure 9, which contains a true positive, false pos-
422 itive, false negative, and true negative case. High values of simulated radar reflectivity
423 (>65) are evident near the storm core of the true positive case (Fig. 9a), which repre-
424 sents a region of high precipitation intensity. The false positive and false negative ex-
425 amples also contain storms with high reflectivity (>65 ; Fig. 9b,c), but the most intense
426 region for the false negative case is located near the southern edge of the image. The true
427 negative case (Fig. 9d) contains convection that is smaller in size, less organized, and
428 of lower maximum reflectivity magnitude than the other examples (<65), which possi-
429 bly contributed to the true negative classification by the CNN.

430 Saliency maps highlight the storm patch areas of input features that contributed
431 to the CNN prediction. For water vapor mixing ratio (right two columns in Fig. 10), pos-
432 itive gradients demarcate the respective pixels that contributed positively to the model
433 prediction. Moisture at low and mid level heights for the true positive case located near
434 the storm core contributed positively to the prediction of strongly rotating storms (Fig.
435 10c,d). While high moisture content may not be related to storm rotation and horizon-
436 tal kinematics, it does show that the CNN identified the storm core (region of high pre-
437 cipitation intensity, and thus moisture content) as relevant for the strongly rotating pre-
438 diction. Non-salient regions of respective variables are zero gradients and therefore cor-
439 respond to pixels that did not contribute to the model prediction. In the case of zonal
440 and meridional winds (left two columns in Fig. 10), input feature values represent both
441 the magnitude and direction of wind flow. For example, positive zonal wind values are
442 winds pointing from west to east and negative values are winds pointing from east to west.
443 Both directions are indicative of storm rotation if in close proximity to each other. Plots
444 a and b in figure 10 show positive and negative gradients in close proximity to each other
445 and at the storm core location, representative of counter-clockwise rotation. This rota-
446 tion signature contributed to the strongly rotating prediction. Similar gradient patterns
447 are present in the maps of the false positive and false negative examples at storm core
448 locations (Fig. 10e-l). Saliency maps for true negative cases are substantially different
449 (Fig. 10m-p)—gradients are no longer present across a small and focused region near the
450 storm core, but rather across broad areas of the storm patch. Additionally, gradients from
451 zonal and meridional winds generally no longer align to form an organized circulation
452 (Fig. 10m,n).

453 Results generated using saliency maps can be corroborated by visualizing the fre-
454 quency of maximum UH for all storms. Figure 11a shows that the CNN is able to cor-
455 rectly identify strongly rotating storms across a broad range of UH values that exceed
456 $75 \text{ m}^2\text{s}^{-2}$ and is therefore able to capture a variety of storm rotation intensities. Storms
457 that were classified as strongly rotating, but evaluated as false alarms because the cor-
458 responding UH did not exceed $75 \text{ m}^2\text{s}^{-2}$, are heavily skewed towards high UH values (mostly
459 contained UH values that exceeded $40 \text{ m}^2\text{s}^{-2}$; Fig. 11b), which past studies have found
460 to also be representative of supercellular convection (Trapp et al., 2011). Missed clas-
461 sifications of strongly rotating storms generally do not consist of large UH magnitudes

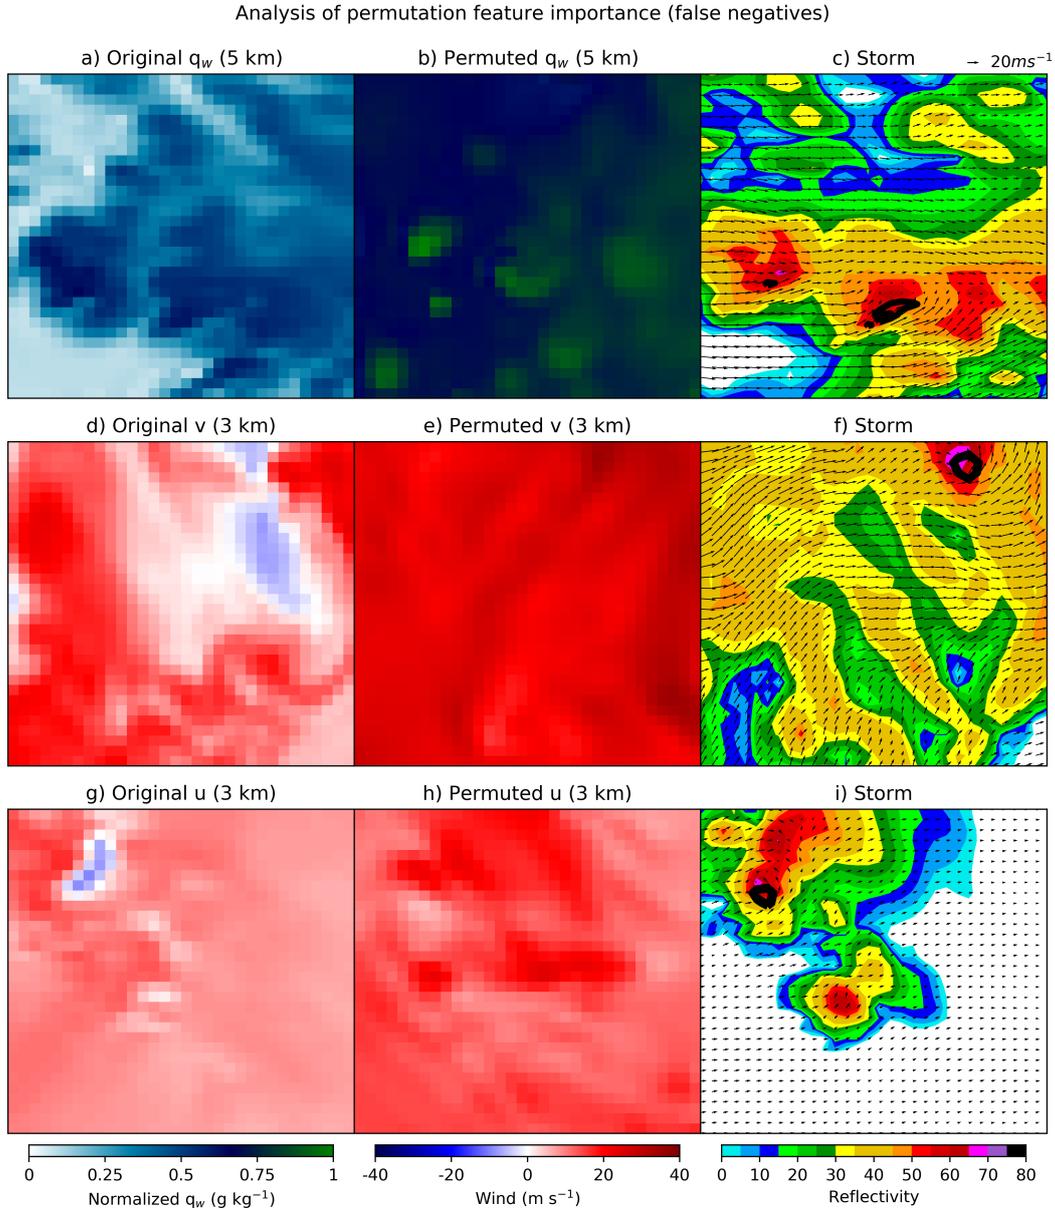


Figure 7. Three example cases (c, f, i) that were incorrectly classified as non-strongly rotating storms during the permutation feature importance (PFI) analysis. The CNN correctly classified these future climate storms as strongly rotating prior to PFI. The top row shows the original water vapor mixing ratio (q_w) field (a) for a pair of strongly-rotating storms (c), and the q_w field that replaced the original during PFI (b). The center and bottom rows show fields for other example storms, but for perturbed meridional (v) and zonal (u) winds respectively. Updraft helicity exceeding $75\ m^2s^{-2}$ is indicated with black contours (c, f, i). Vectors show winds at 3 km corresponding to the respective storm image (c, f, i). The q_w fields were normalized by the maximum value of both plots (a, b).

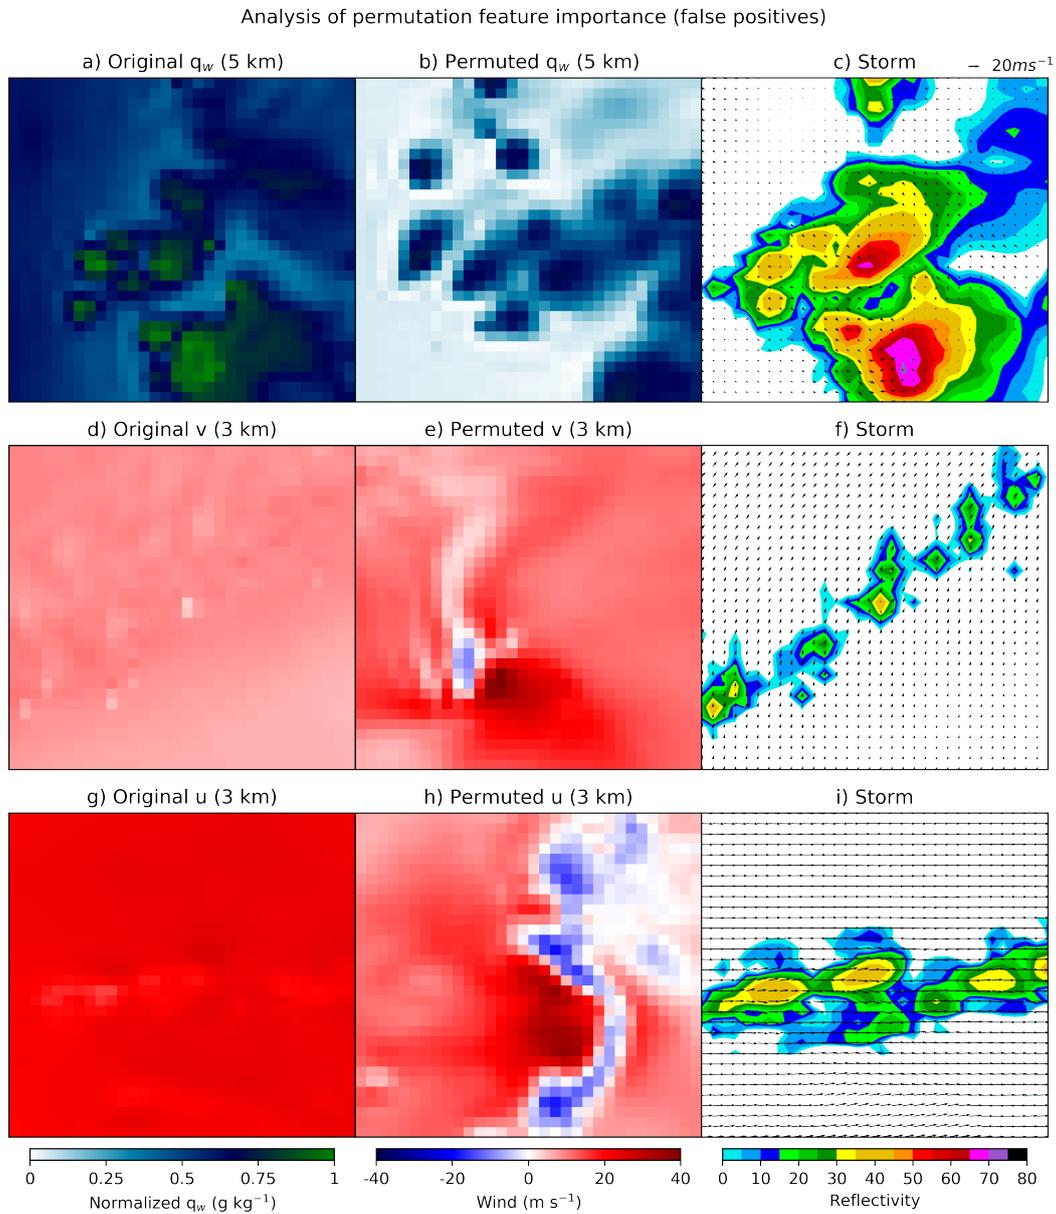


Figure 8. Same as figure 7, but for a subset of storms that were incorrectly classified as strongly rotating. The CNN correctly classified these storms as non-strongly rotating prior to PFI. No black contours are included in the storm plots (c, f, i) because updraft helicity did not exceed $75 \text{ m}^2\text{s}^{-2}$ for the shown examples.

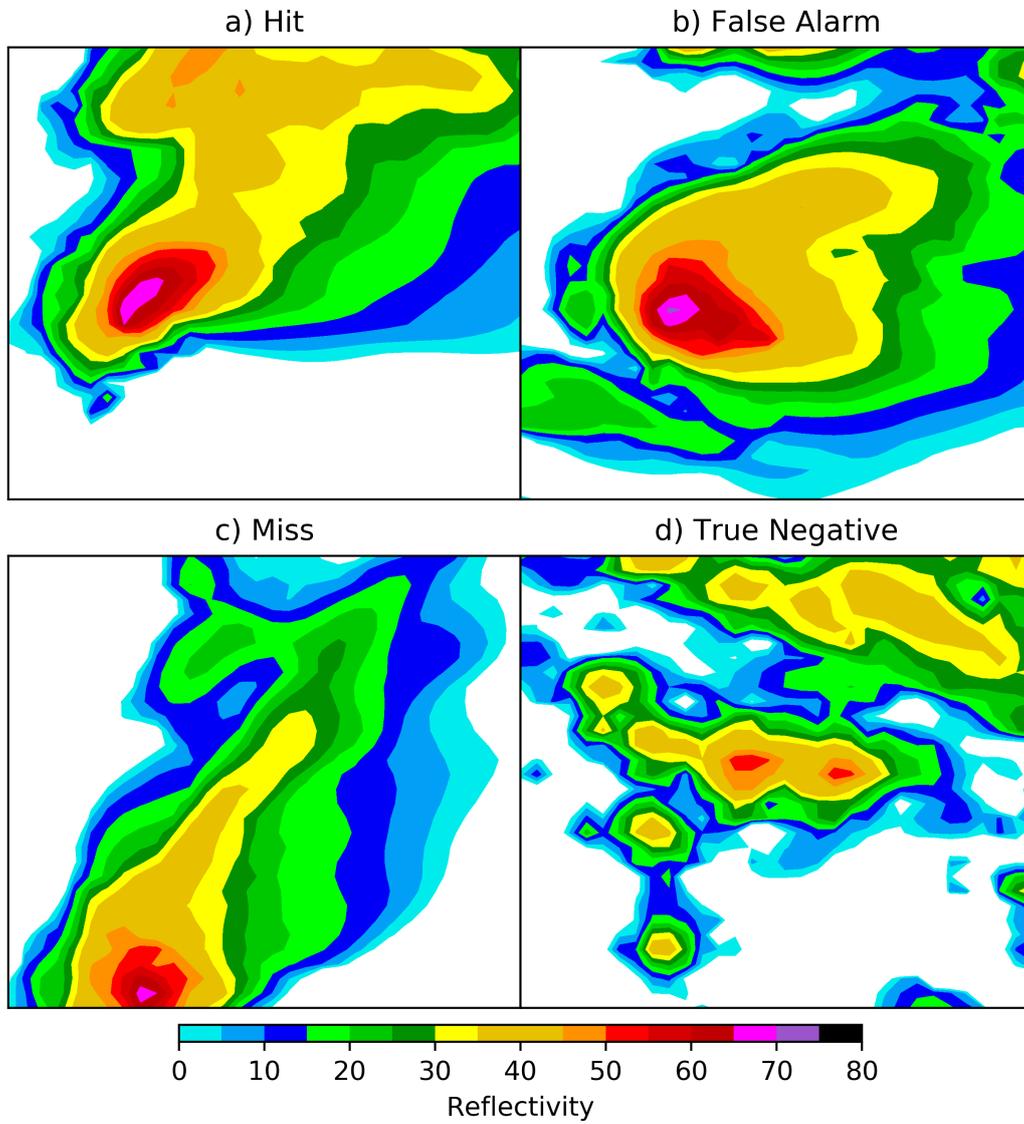


Figure 9. Simulated radar reflectivity for example storms extracted from the future climate simulation, evaluated as a hit (a), false alarm (b), miss (c), and correct negative (d).

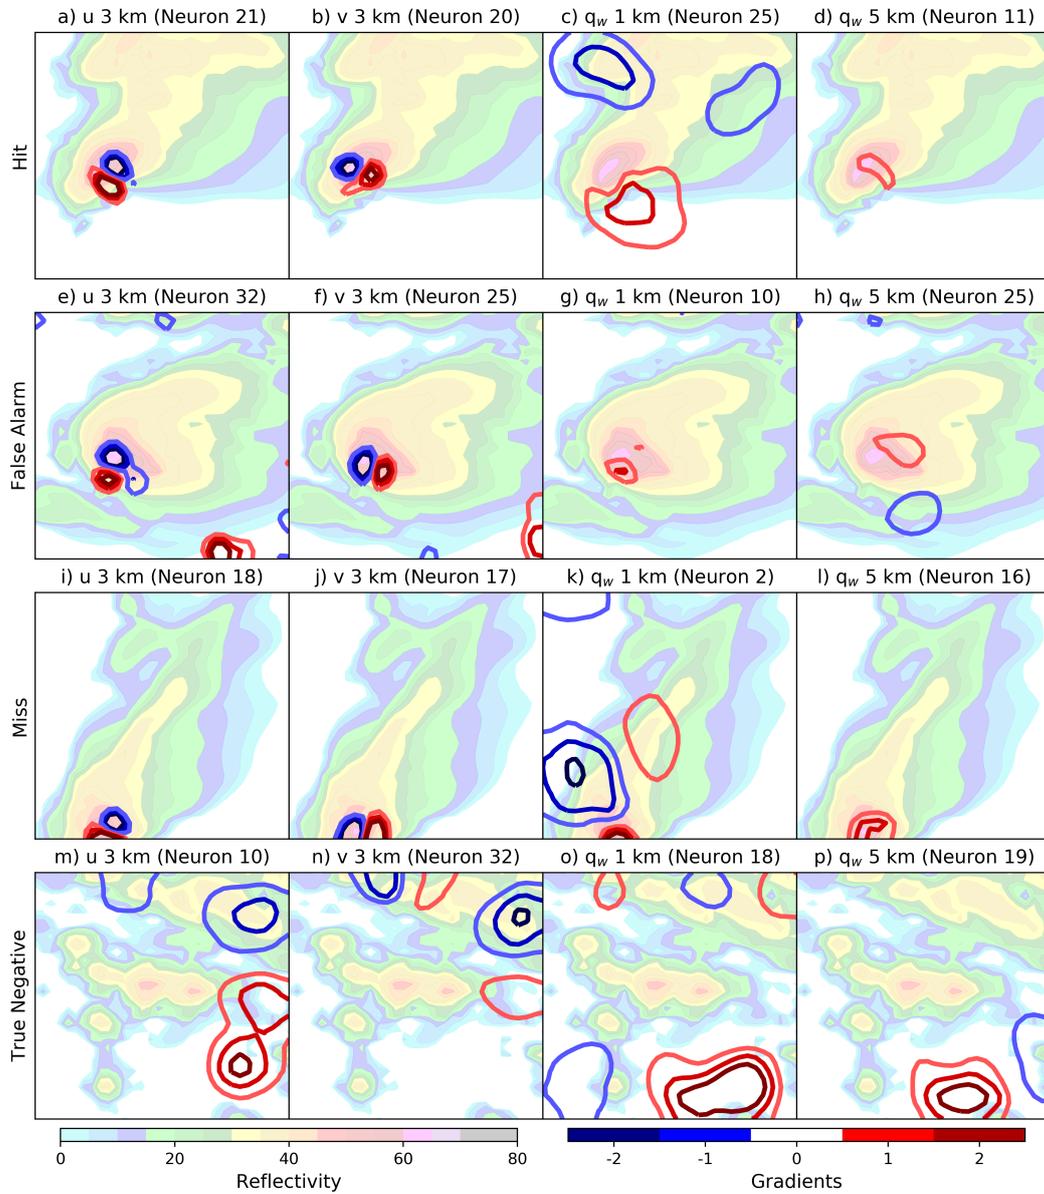


Figure 10. Saliency maps for representative example future climate storms, including true positive (a-d), false positive (e-h), false negative (i-l), and true negative (m-p) cases. Variables shown include several denoted as important by the permutation feature importance analysis, such as 3 km zonal (a,e,i,m) and meridional (b,f,j,n) winds, and water vapor mixing ratio (q_w) at 5 km (d,h,l,p). Mixing ratio (q_w) at 1 km is also shown (c,g,k,o).

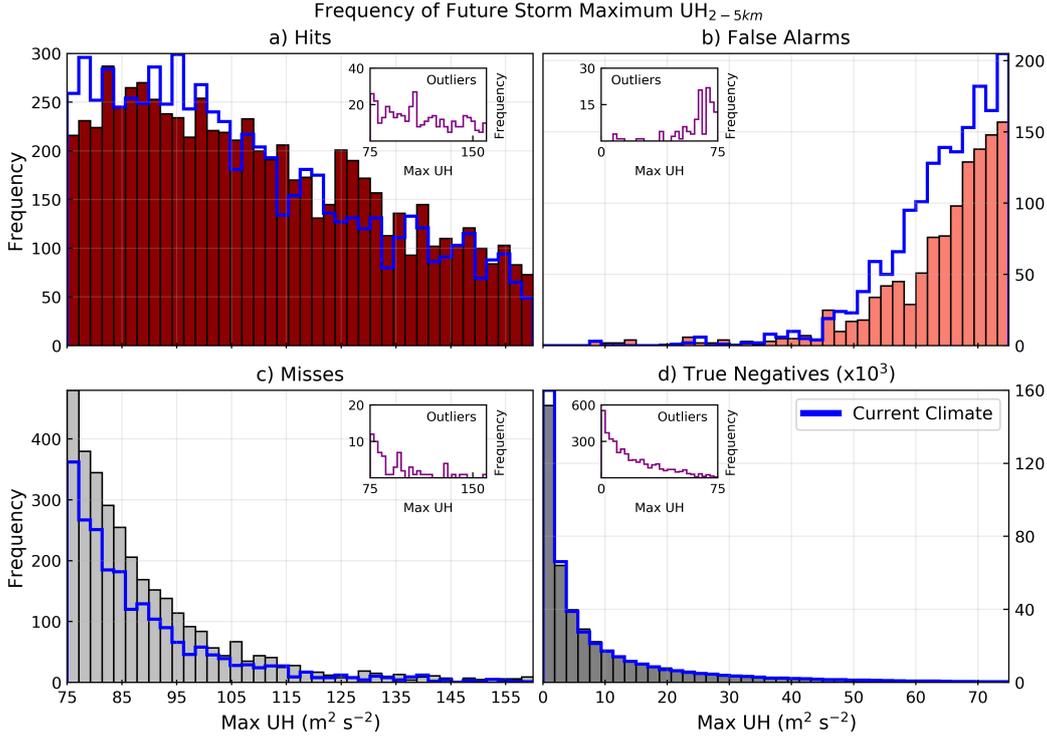


Figure 11. Histograms show the frequency of maximum updraft helicity (UH) for future climate storms separated into four subsets: hits (a), false alarms (b), misses (c), and true negatives (d). Frequency of true negatives (d) are $\times 10^3$ magnitude. For comparison, the frequencies of the maximum UH for current climate storms are shown with blue lines and for future outlier storms in the inset plots.

462 ($< 100 \text{ m}^2 \text{ s}^{-2}$), with most storms characterized by UH values close to $75 \text{ m}^2 \text{ s}^{-2}$ (Fig. 11c).
 463 Most true negative cases consist of UH values below $40 \text{ m}^2 \text{ s}^{-2}$, which is characteristic
 464 of less organized convective storms (Fig. 11d). These results further demonstrate that
 465 a CNN is able to generalize a target derived from a heuristic using learned features in
 466 the data that would be difficult to encode due to spatial complexity. There is sensitiv-
 467 ity, however, to the storm location within the storm patch, which can be visualized with
 468 2D histograms that contain the frequency of UH exceeding $75 \text{ m}^2 \text{ s}^{-2}$, typically located
 469 near the thunderstorm core (Fig. 12). Correctly classified storms contained regions of
 470 high rotation ($\text{UH} > 75 \text{ m}^2 \text{ s}^{-2}$) near the center of the storm patch (Fig. 12a,c), while missed
 471 classifications are located near the edges of the storm patch (Fig. 12b,d) for storms dur-
 472 ing both the current and future climate. This comparison shows that CNNs can strug-
 473 gle with correct classifications of features located near the edges of a spatial region of
 474 interest, resulting in missed events.

475 4 Conclusions

476 A CNN was trained to learn relationships and identify features among meteorological
 477 state variables in order to classify convection types, with a focus on rotation within
 478 the updraft core of a thunderstorm. Strong rotation and associated storm morphology
 479 could result in a higher likelihood of convection producing severe hazards, such as tor-
 480 nadoes and large hail, which are a dangerous threat to the public. We hypothesized that
 481 due to climate change, a trained CNN may fail to classify and identify convection that

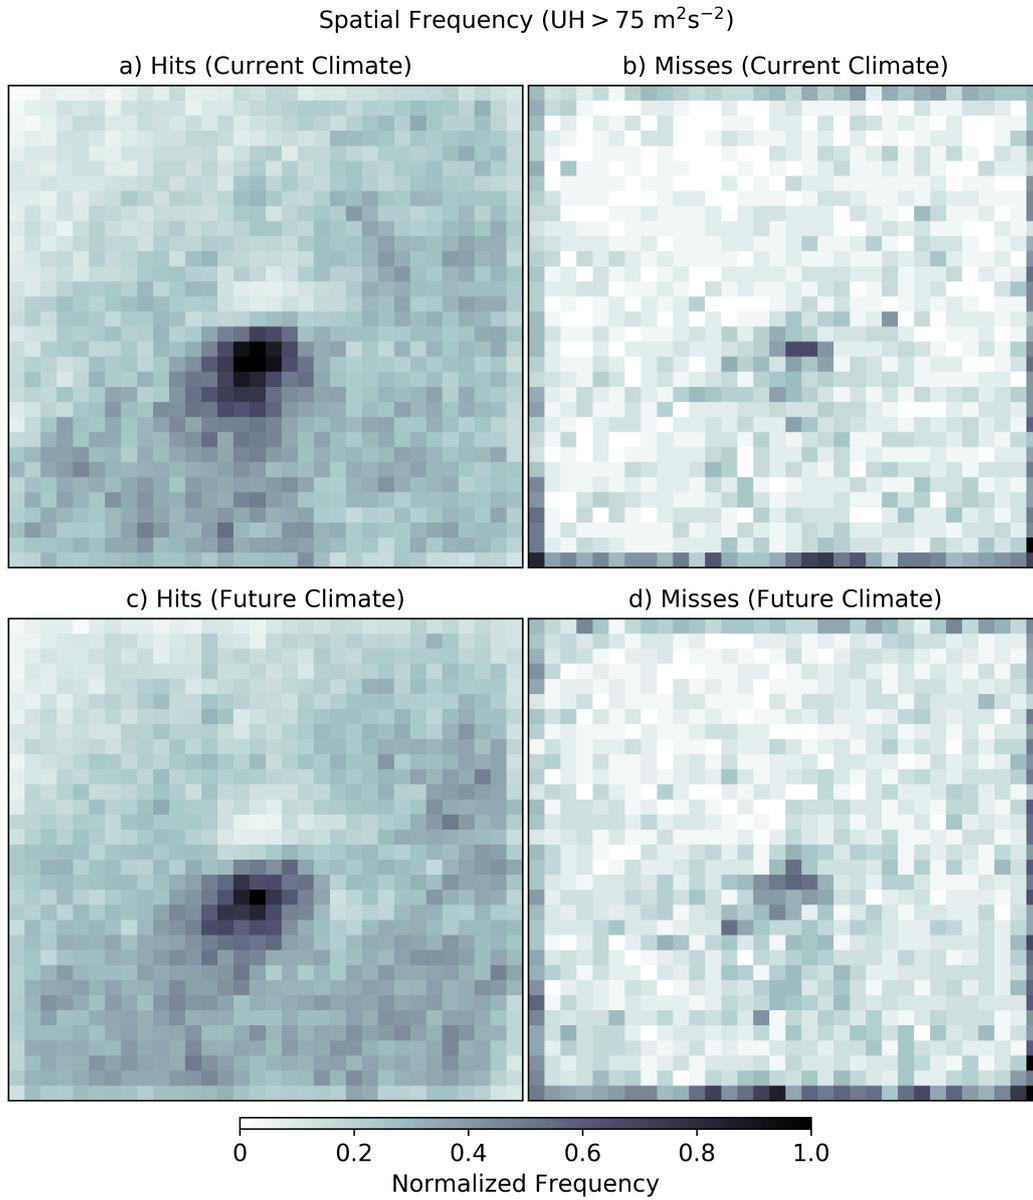


Figure 12. Spatial histograms show the frequency of updraft helicity (UH) exceeding $75 \text{ m}^2\text{s}^{-2}$ normalized by maximum frequency for storm objects classified as hits (a,c) and misses (b,d) during the current (a,b) and future climates (c,d).

482 lies outside of the climatological distribution of data used for training. Using a thermo-
483 dynamically driven future climate model simulation, we show that a CNN can remain
484 skillful in classifying rotating convective storms via learned representations of physical
485 variables.

486 The key results that provide answers to the questions posed in the introduction follow:

- 487 1. A CNN trained using a current climate model simulation can skillfully classify out-
488 of-sample (with regards to moisture content) storms in a thermodynamically driven
489 future climate. This is likely partly due to the use of batch normalization and spa-
490 tial dropout; an equivalent model trained without batch normalization and spa-
491 tial dropout results in a more substantial under-forecasting bias (about 0.75) across
492 climate subgroups.
- 493 2. Kinematic fields and mid-level moisture were identified as important variables for
494 skillful classification by the CNN. Spatially, wind rotation signatures with concur-
495 rently overlaid sharp mid-level moisture gradients were also important.
- 496 3. Incorrect storm classifications included cases that were near the storm patch edge
497 or had a UH value that was near but on the opposite side of the predefined thresh-
498 old.

499 Key result 1 shows that a CNN is robust to out-of-sample cases during convection
500 classification, which is a promising result given the changes already occurring to large-
501 scale environments and moisture advection patterns associated with severe thunderstorms
502 (Gensini & Brooks, 2018; Molina & Allen, 2020). Key results 2 and 3 also show that a
503 CNN can learn complex relationships among input features using labels derived from heuris-
504 tics. Physical features were not prescribed but rather learned from the data, such as the
505 importance of dry air at mid-levels for intense storm development when low-level mois-
506 ture is present (i.e., convective available potential energy). Unlike computer vision clas-
507 sification tasks (Russakovsky et al., 2015), humans can bypass generating a large num-
508 ber of hand labeled data for training models to perform atmospheric feature classifica-
509 tions, which would also pose challenges given conflicting definitions of atmospheric phe-
510 nomena in the scientific literature. Additionally, results show that large imbalances in
511 labeled data may be overcome with sufficient hyperparameter tuning. Overall, results
512 show that the CNN can classify storms as strongly rotating that were near the UH thresh-
513 old and appeared supercellular, learning to generalize prescribed UH labels.

514 There are several limitations that are important to acknowledge, however. The fo-
515 cus in this study lies on a future climate that was thermodynamically driven in order
516 to isolate competing thermodynamic and kinematic signals, but it is possible that a CNN
517 may not generalize well with a future climate that accounts for both changes in the ther-
518 modynamic and large-scale dynamics. We do note that there is a 14% increase in future
519 strongly rotating storms as compared to the current climate, which is a substantial in-
520 crease and an indication that changes in large-scale dynamics may not pose a significant
521 issue to the CNN. An additional limitation is that physical interpretation methods re-
522 quire substantial human interpretation, making it possible to miss important features
523 or fail to discover new physical relationships. However, this is a broader issue within ma-
524 chine learning interpretability, as it introduces the potential for confirmation bias from
525 human scientists attempting to explain results. Future work should explore incorporat-
526 ing feature uncertainty or physics within the CNN model architecture to explore the dif-
527 ferences to results contained herein. Additionally, methods to ameliorate missed clas-
528 sifications near the edges of study domains should be explored. As societal exposure to
529 severe hazards continues to increase (e.g., Ashley & Strader, 2016), it is important to
530 continue better identifying and understanding severe hazards within climate model sim-
531 ulations. The use of deep learning methods that do not impose rigid thresholds or ex-
532 pert systems decisions should continue to be explored, since meteorological phenomena

generally do not neatly fit into predefined classes. Deep learning offers a viable avenue to continue to better understand weather and climate extremes.

Acknowledgments

This research was funded by the Advanced Study Program Postdoctoral Fellowship of the National Center for Atmospheric Research (NCAR), which is sponsored by the National Science Foundation. High-performance computing support on Cheyenne (doi:10.5065/D6RX99HX) was provided by the NCAR Computational and Information Systems Laboratory (CISL). The climate data set (doi:10.5065/D6V40SXP R. Rasmussen & Liu, 2017) can be freely obtained from the NCAR Research Data Archive hosted by CISL at <https://rda.ucar.edu/datasets/ds612.0/>. The software developed for this research is available on Github <https://github.com/mariajmolina/deep-conus>. Processed data shown in the figures is archived in Zenodo (doi:10.5281/zenodo.4052586 Molina, Gagne, & Prein, 2020). The convolutional neural network was trained using the high-level Python-based Keras deep learning library (Chollet et al., 2015) with the TensorFlow low-level backend (Abadi et al., 2016).

References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., . . . others (2016). Tensorflow: A system for large-scale machine learning. In *12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16)* (pp. 265–283).
- Allen, J. T. (2018). Climate change and severe thunderstorms. *Oxford Research Encyclopedia of Climate Science*. doi: 10.1093/acrefore/9780190228620.013.62
- Ashley, W. S., & Strader, S. M. (2016). Recipe for disaster: How the dynamic ingredients of risk and exposure are changing the tornado disaster landscape. *Bull. Am. Meteor. Soc.*, *97*(5), 767–786. doi: 10.1175/BAMS-D-15-00150.1
- Barlage, M., Chen, F., Miguez-Macho, G., Liu, C., Liu, X., & Niyogi, D. (2018). Enhancing hydrologic processes in the Noah-MP land surface model to improve seasonal forecast skill. In *98th american meteorological society annual meeting*.
- Barnes, E. A., Hurrell, J. W., Ebert-Uphoff, I., Anderson, C., & Anderson, D. (2019). Viewing forced climate patterns through an AI lens. *Geophys. Res. Lett.*, *46*(22), 13389–13398. doi: 10.1029/2019GL084944
- Biard, J. C., & Kunkel, K. E. (2019). Automated detection of weather fronts using a deep learning neural network. *Advances in Statistical Climatology, Meteorology and Oceanography*, *5*(2), 147–160. doi: 10.5194/ascmo-5-147-2019
- Breiman, L. (2001). Random forests. *Machine learning*, *45*(1), 5–32. doi: 10.1023/a:1010933404324
- Brooks, H. E. (2013). Severe thunderstorms and climate change. *Atmos. Res.*, *123*, 129–138. doi: 10.1016/j.atmosres.2012.04.002
- Brooks, H. E., Doswell III, C. A., & Kay, M. P. (2003). Climatological estimates of local daily tornado probability for the United States. *Wea. Forecasting*, *18*(4), 626–640. doi: 10.1175/1520-0434(2003)018<0626:CEOLDT>2.0.CO;2
- Bunkers, M. J., Hjelmfelt, M. R., & Smith, P. L. (2006). An observational examination of long-lived supercells. Part I: Characteristics, evolution, and demise. *Wea. Forecasting*, *21*(5), 673–688. doi: 10.1175/WAF949.1
- Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: synthetic minority over-sampling technique. *J. Artificial Intelligence Res.*, *16*, 321–357. doi: 10.1613/jair.953
- Chollet, F., et al. (2015). Keras documentation. *keras.io*.
- Clark, A. J., Gao, J., Marsh, P. T., Smith, T., Kain, J. S., Correia Jr, J., . . . Kong, F. (2013). Tornado pathlength forecasts from 2010 to 2011 using ensemble updraft helicity. *Wea. Forecasting*, *28*(2), 387–407. doi:

- 584 10.1175/WAF-D-12-00038.1
- 585 Dee, D. P., Uppala, S. M., Simmons, A. J., Berrisford, P., Poli, P., Kobayashi, S.,
586 ... Vitart, F. (2011). The ERA-Interim reanalysis: Configuration and perfor-
587 mance of the data assimilation system. *Quart. J. Roy. Meteor. Soc.*, *137*(656),
588 553–597. doi: 10.1002/qj.828
- 589 Diffenbaugh, N. S., Scherer, M., & Trapp, R. J. (2013). Robust increases in se-
590 vere thunderstorm environments in response to greenhouse forcing. *Proc. Natl.*
591 *Acad. Sci. (USA)*, *110*(41), 16361–16366. doi: 10.1073/pnas.1307758110
- 592 Doswell, C. A., III, Brooks, H. E., & Maddox, R. A. (1996). Flash flood forecasting:
593 An ingredients-based methodology. *Wea. Forecasting*, *11*(4), 560–581. doi: 10
594 .1175/1520-0434(1996)011<0560:FFFAIB>2.0.CO;2
- 595 Duda, J. D., & Gallus Jr, W. A. (2010). Spring and summer midwestern severe
596 weather reports in supercells compared to other morphologies. *Wea. Forecast-*
597 *ing*, *25*(1), 190–206. doi: 10.1175/2009WAF2222338.1
- 598 Fukushima, K., & Miyake, S. (1982). Neocognitron: A new algorithm for pattern
599 recognition tolerant of deformations and shifts in position. *Pattern recognition*,
600 *15*(6), 455–469. doi: 10.1007/978-3-642-46466-9_18
- 601 Gagne II, D. J., Haupt, S. E., Nychka, D. W., & Thompson, G. (2019).
602 Interpretable deep learning for spatial analysis of severe hailstorms.
603 *Mon. Wea. Rev.*, *147*(8), 2827–2845. doi: 10.1175/MWR-D-18-0316.1
- 604 Gensini, V. A., & Brooks, H. E. (2018). Spatial trends in United States tornado fre-
605 quency. *npj Climate and Atmospheric Science*, *1*, 38. doi: 10.1038/s41612-018
606 -0048-2
- 607 Gensini, V. A., & Mote, T. L. (2015). Downscaled estimates of late 21st century se-
608 vere weather from CCSM3. *Climate Change*, *129*(1–2), 307–321. doi: 10.1007/
609 s10584-014-1320-z
- 610 Goodfellow, I., Bengio, Y., Courville, A., & Bengio, Y. (2016). *Deep learning*
611 (Vol. 1). MIT press Cambridge.
- 612 Gropp, M. E., & Davenport, C. E. (2018). The impact of the nocturnal transition on
613 the lifetime and evolution of supercell thunderstorms in the Great Plains. *Wea.*
614 *Forecasting*, *33*(4), 1045–1061. doi: 10.1175/WAF-D-17-0150.1
- 615 Ham, Y.-G., Kim, J.-H., & Luo, J.-J. (2019). Deep learning for multi-year ENSO
616 forecasts. *Nature*, *573*(7775), 568–572. doi: s41586-019-1559-7
- 617 Hong, S.-Y., Noh, Y., & Dudhia, J. (2006). A new vertical diffusion package with
618 an explicit treatment of entrainment processes. *Mon. Wea. Rev.*, *134*(9), 2318–
619 2341. doi: 10.1175/MWR3199.1
- 620 Hoogewind, K. A., Baldwin, M. E., & Trapp, R. J. (2017). The impact of climate
621 change on hazardous convective weather in the United States: Insight from
622 high-resolution dynamical downscaling. *J. Climate*, *30*(24), 10081–10100. doi:
623 10.1175/JCLI-D-16-0885.1
- 624 Hsu, W.-r., & Murphy, A. H. (1986). The attributes diagram a geometrical frame-
625 work for assessing the quality of probability forecasts. *International Journal of*
626 *Forecasting*, *2*(3), 285–293. doi: 10.1016/0169-2070(86)90048-8
- 627 Iacono, M. J., Delamere, J. S., Mlawer, E. J., Shephard, M. W., Clough, S. A., &
628 Collins, W. D. (2008). Radiative forcing by long-lived greenhouse gases: Cal-
629 culations with the AER radiative transfer models. *J. Geophys. Res.: Atmos.*,
630 *113*(D13). doi: 10.1029/2008JD009944
- 631 Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network
632 training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*.
- 633 Jergensen, G. E., McGovern, A., Lagerquist, R., & Smith, T. (2020). Classifying
634 convective storms using machine learning. *Wea. Forecasting*, *35*(2), 537–559.
635 doi: 10.1175/WAF-D-19-0170.1
- 636 Kain, J. S., Weiss, S. J., Bright, D. R., Baldwin, M. E., Levit, J. J., Carbin, G. W.,
637 ... Thomas, K. (2008). Some practical considerations regarding horizontal res-
638 olution in the first generation of operational convection-allowing NWP. *Wea.*

- 639 *Forecasting*, 23(5), 931–952. doi: 10.1175/WAF2007106.1
- 640 Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv*
- 641 *preprint arXiv:1412.6980*.
- 642 Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with
- 643 deep convolutional neural networks. In *Advances in neural information pro-*
- 644 *cessing systems* (pp. 1097–1105). Retrieved from [http://papers.nips.cc/](http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf)
- 645 [paper/4824-imagenet-classification-with-deep-convolutional-neural](http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf)
- 646 [-networks.pdf](http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf)
- 647 Lagerquist, R., McGovern, A., & Gagne II, D. J. (2019). Deep learning for spa-
- 648 tially explicit prediction of synoptic-scale fronts. *Wea. Forecasting*, 34(4),
- 649 1137–1160. doi: 10.1175/WAF-D-18-0183.1
- 650 Lagerquist, R., McGovern, A., Homeyer, C. R., Gagne, D. J., & Smith, T. (2020).
- 651 Deep learning on three-dimensional multiscale data for next-hour tornado
- 652 prediction. *Mon. Wea. Rev.*, 148(7). doi: 10.1175/MWR-D-19-0372.1
- 653 Lakshmanan, V., Hondl, K., & Rabin, R. (2009). An efficient, general-purpose tech-
- 654 nique for identifying storm cells in geospatial images. *Journal of Atmospheric*
- 655 *and Oceanic Technology*, 26(3), 523–537. doi: 10.1175/2008JTECHA1153.1
- 656 Lakshmanan, V., Karstens, C., Krause, J., Elmore, K., Ryzhkov, A., & Berkseth, S.
- 657 (2015). Which polarimetric variables are important for weather/no-weather
- 658 discrimination? *Journal of Atmospheric and Oceanic Technology*, 32(6),
- 659 1209–1223. doi: 10.1175/JTECH-D-13-00205.1
- 660 LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553),
- 661 436–444. doi: 10.1038/nature14539
- 662 LeCun, Y., Boser, B. E., Denker, J. S., Henderson, D., Howard, R. E., Hubbard,
- 663 W. E., & Jackel, L. D. (1990). Handwritten digit recognition with a back-
- 664 propagation network. In *Advances in neural information processing systems*
- 665 (pp. 396–404).
- 666 LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning
- 667 applied to document recognition. *Proceedings of the IEEE*, 86(11), 2278–2324.
- 668 doi: 10.1109/5.726791
- 669 Liu, C., Ikeda, K., Rasmussen, R., Barlage, M., Newman, A. J., Prein, A. F., ...
- 670 Dudhia, J. (2017). Continental-scale convection-permitting modeling of the
- 671 current and future climate of North America. *Clim. Dyn.*, 49(1–2), 71–95. doi:
- 672 10.1007/s00382-016-3327-9
- 673 Liu, Y., Racah, E., Correa, J., Khosrowshahi, A., Lavers, D., Kunkel, K., ...
- 674 Collins, W. (2016). Application of deep convolutional neural networks for de-
- 675 tecting extreme weather in climate datasets. *arXiv preprint arXiv:1605.01156*.
- 676 Mason, I. (1982). A model for assessment of weather forecasts. *Aust. Meteor. Mag.*,
- 677 30(4), 291–303.
- 678 McGovern, A., Lagerquist, R., John Gagne, D., Jergensen, G. E., Elmore, K. L.,
- 679 Homeyer, C. R., & Smith, T. (2019). Making the black box more transparent:
- 680 Understanding the physical implications of machine learning. *Bull. Am. Me-*
- 681 *teor. Soc.*, 100(11), 2175–2199. doi: 10.1175/BAMS-D-18-0195.1
- 682 Molina, M. J., & Allen, J. T. (2020). Regionally-stratified tornadoes: Moisture
- 683 source physical reasoning and climate trends. *Wea. and Clim. Extremes*, 28,
- 684 100244. doi: 10.1016/j.wace.2020.100244
- 685 Molina, M. J., Allen, J. T., & Prein, A. F. (2020). Moisture attribution and sensitiv-
- 686 ity analysis of a winter tornado outbreak. *Wea. Forecasting*, 35(4), 1263–1288.
- 687 doi: 10.1175/WAF-D-19-0240.1
- 688 Molina, M. J., Gagne, D. J., & Prein, A. F. (2020). *DATA: Deep learning classifi-*
- 689 *cation of potentially severe convective storms in a changing climate [Data set]*.
- 690 Zenodo. Retrieved from <https://doi.org/10.5281/zenodo.4052586>
- 691 Moss, R. H., Edmonds, J. A., Hibbard, K. A., Manning, M. R., Rose, S. K., Van Vu-
- 692 uren, D. P., ... Meehl, G. A. (2010). The next generation of scenarios for
- 693 climate change research and assessment. *Nature*, 463(7282), 747–756. doi:

- 694 10.1038/nature08823
- 695 Niu, G.-Y., Yang, Z.-L., Mitchell, K. E., Chen, F., Ek, M. B., Barlage, M., ...
 696 Tewari, M. (2011). The community Noah land surface model with multi-
 697 parameterization options (Noah-MP): 1. Model description and evaluation
 698 with local-scale measurements. *J. Geophys. Res.: Atmos.*, *116*(D12). doi:
 699 10.1029/2010JD015139
- 700 Prein, A. F., Liu, C., Ikeda, K., Trier, S. B., Rasmussen, R. M., Holland, G. J.,
 701 & Clark, M. P. (2017). Increased rainfall volume from future convective
 702 storms in the us. *Nature Climate Change*, *7*(12), 880–884. doi:
 703 10.1038/s41558-017-0007-7
- 704 Rasmussen, E. N. (2003). Refined supercell and tornado forecast parameters. *Wea.*
 705 *Forecasting*, *18*(3), 530–535. doi: 10.1175/1520-0434(2003)18(530:RSATFP)2.0
 706 .CO;2
- 707 Rasmussen, K. L., Prein, A. F., Rasmussen, R. M., Ikeda, K., & Liu, C. (2017).
 708 Changes in the convective population and thermodynamic environments in
 709 convection-permitting regional climate simulations over the united states.
 710 *Climate Dynamics*, *55*, 1–26. doi: 10.1007/s00382-017-4000-7
- 711 Rasmussen, R., & Liu, C. (2017). *High resolution WRF simulations of the current*
 712 *and future climate of North America*. Research Data Archive at the National
 713 Center for Atmospheric Research, Computational and Information Systems
 714 Laboratory. Retrieved from <https://doi.org/10.5065/D6V40SXP>
- 715 Rasp, S., Pritchard, M. S., & Gentine, P. (2018). Deep learning to represent subgrid
 716 processes in climate models. *Proceedings of the National Academy of Sciences*,
 717 *115*(39), 9684–9689. doi: 10.1073/pnas.1810286115
- 718 Roebber, P. J. (2009). Visualizing multiple measures of forecast quality. *Wea. Fore-*
 719 *casting*, *24*(2), 601–608. doi: 10.1175/2008WAF2222159.1
- 720 Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., ... Fei-Fei,
 721 L. (2015). ImageNet Large Scale Visual Recognition Challenge. *In-*
 722 *ternational Journal of Computer Vision (IJCV)*, *115*(3), 211-252. doi:
 723 10.1007/s11263-015-0816-y
- 724 Schär, C., Frei, C., Lüthi, D., & Davies, H. C. (1996). Surrogate climate-change sce-
 725 narios for regional climate models. *Geophys. Res. Lett.*, *23*(6), 669–672. doi: 10
 726 .1029/96GL00265
- 727 Simonyan, K., Vedaldi, A., & Zisserman, A. (2013). Deep inside convolutional
 728 networks: Visualising image classification models and saliency maps. *arXiv*
 729 *preprint arXiv:1312.6034*.
- 730 Skamarock, W. C., & Klemp, J. B. (2008). A time-split nonhydrostatic atmospheric
 731 model for weather research and forecasting applications. *Journal of computa-*
 732 *tional physics*, *227*(7), 3465–3485. doi: 10.1016/j.jcp.2007.01.037
- 733 Sobash, R. A., & Kain, J. S. (2017). Seasonal variations in severe weather forecast
 734 skill in an experimental convection-allowing model. *Wea. Forecasting*, *32*(5),
 735 1885–1902. doi: 10.1175/WAF-D-17-0043.1
- 736 Sobash, R. A., Kain, J. S., Bright, D. R., Dean, A. R., Coniglio, M. C., & Weiss,
 737 S. J. (2011). Probabilistic forecast guidance for severe thunderstorms based
 738 on the identification of extreme phenomena in convection-allowing model fore-
 739 casts. *Wea. Forecasting*, *26*(5), 714–728. doi: 10.1175/WAF-D-10-05046.1
- 740 Sobash, R. A., Romine, G. S., & Schwartz, C. S. (2020). A comparison of neural-
 741 network and surrogate-severe probabilistic convective hazard guidance derived
 742 from a convection-allowing model. *Wea. Forecasting*, *35*(5), 1981–2000. doi:
 743 10.1175/WAF-D-20-0036.1
- 744 Sobash, R. A., Schwartz, C. S., Romine, G. S., Fossell, K. R., & Weisman,
 745 M. L. (2016). Severe weather prediction using storm surrogates from
 746 an ensemble forecasting system. *Wea. Forecasting*, *31*(1), 255–271. doi:
 747 10.1175/WAF-D-15-0138.1
- 748 Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R.

- 749 (2014). Dropout: a simple way to prevent neural networks from overfit-
750 ting. *The journal of machine learning research*, *15*(1), 1929–1958. doi:
751 10.5555/2627435.2670313
- 752 Taszarek, M., Allen, J. T., Brooks, H. E., Pilguy, N., & Czernecki, B. (2020).
753 Differing trends in United States and European severe thunderstorm en-
754 vironments in a warming climate. *Bull. Am. Meteor. Soc.*, 1–51. doi:
755 10.1175/BAMS-D-20-0004.1
- 756 Taylor, K. E., Stouffer, R. J., & Meehl, G. A. (2012). An overview of CMIP5 and
757 the experiment design. *Bull. Am. Meteor. Soc.*, *93*(4), 485–498. doi: 10.1175/
758 BAMS-D-11-00094.1
- 759 Thompson, G., & Eidhammer, T. (2014). A study of aerosol impacts on clouds and
760 precipitation development in a large winter cyclone. *J. Atmos. Sci.*, *71*(10),
761 3636–3658. doi: 10.1175/JAS-D-13-0305.1
- 762 Thompson, R. L., Edwards, R., Hart, J. A., Elmore, K. L., & Markowski, P.
763 (2003). Close proximity soundings within supercell environments obtained
764 from the Rapid Update Cycle. *Wea. Forecasting*, *18*(6), 1243–1261. doi:
765 10.1175/1520-0434(2003)018<1243:CPSWSE>2.0.CO;2
- 766 Thompson, R. L., Mead, C. M., & Edwards, R. (2007). Effective storm-relative he-
767 licity and bulk shear in supercell thunderstorm environments. *Wea. Forecast-*
768 *ing*, *22*(1), 102–115. doi: 10.1175/WAF969.1
- 769 Thompson, R. L., Smith, B. T., Grams, J. S., Dean, A. R., & Broyles, C. (2012).
770 Convective modes for significant severe thunderstorms in the contiguous
771 United States. Part II: Supercell and QLCS tornado environments. *Wea.*
772 *Forecasting*, *27*(5), 1136–1154. doi: 10.1175/WAF-D-11-00116.1
- 773 Toms, B. A., Kashinath, K., Yang, D., et al. (2019). Deep learning for scientific
774 inference from geophysical data: The Madden-Julian Oscillation as a test case.
775 *arXiv preprint arXiv:1902.04621*.
- 776 Trapp, R. J., Diffsenbaugh, N. S., Brooks, H. E., Baldwin, M. E., Robinson, E. D.,
777 & Pal, J. S. (2007). Changes in severe thunderstorm environment frequency
778 during the 21st century caused by anthropogenically enhanced global ra-
779 diative forcing. *Proc. Natl. Acad. Sci. (USA)*, *104*(50), 19719–19723. doi:
780 10.1073/pnas.0705494104
- 781 Trapp, R. J., Diffsenbaugh, N. S., & Gluhovsky, A. (2009). Transient response of se-
782 vere thunderstorm forcing to elevated greenhouse gas concentrations. *Geophys.*
783 *Res. Lett.*, *36*(1). doi: 10.1029/2008GL036203
- 784 Trapp, R. J., & Hoogewind, K. A. (2016). The realization of extreme tornadic storm
785 events under future anthropogenic climate change. *J. Climate*, *29*(14), 5251–
786 5265. doi: 10.1175/JCLI-D-15-0623.1
- 787 Trapp, R. J., Robinson, E. D., Baldwin, M. E., Diffsenbaugh, N. S., & Schwedler,
788 B. R. (2011). Regional climate of hazardous convective weather through
789 high-resolution dynamical downscaling. *Clim. Dyn.*, *37*(3–4), 677–688. doi:
790 10.1007/s00382-010-0826-y
- 791 von Storch, H., Langenberg, H., & Feser, F. (2000). A spectral nudging technique for
792 dynamical downscaling purposes. *Mon. Wea. Rev.*, *128*(10), 3664–3673. doi:
793 10.1175/1520-0493%282000%29128<3664%3AASNTFD>2.0.CO%3B2
- 794 Wandishin, M. S., Baldwin, M. E., Mullen, S. L., & Cortinas Jr, J. V. (2005). Short-
795 range ensemble forecasts of precipitation type. *Wea. Forecasting*, *20*(4), 609–
796 626. doi: 10.1175/WAF871.1
- 797 Wilks, D. S. (2011). *Statistical methods in the atmospheric sciences* (Vol. 100). Aca-
798 demic press.
- 799 Zhou, K., Zheng, Y., Li, B., Dong, W., & Zhang, X. (2019). Forecasting different
800 types of convective weather: A deep learning approach. *Journal of Meteorolog-*
801 *ical Research*, *33*(5), 797–809. doi: 10.1007/s13351-019-8162-6