Climatology and Variability of Warm and Cold Fronts over North America from 1979 to 2018

RYAN LAGERQUIST

School of Meteorology, University of Oklahoma, Norman, Oklahoma

JOHN T. ALLEN

Department of Earth and Atmospheric Sciences, Central Michigan University, Mt. Pleasant, Michigan

AMY MCGOVERN

School of Computer Science, University of Oklahoma, Norman, Oklahoma

(Manuscript received 14 September 2019, in final form 24 April 2020)

ABSTRACT

This paper describes the development and analysis of an objective climatology of warm and cold fronts over North America from 1979 to 2018. Fronts are detected by a convolutional neural network (CNN), trained to emulate fronts drawn by human meteorologists. Predictors for the CNN are surface and 850-hPa fields of temperature, specific humidity, and vector wind from the ERA5 reanalysis. Gridded probabilities from the CNN are converted to 2D frontal regions, which are used to create the climatology. Overall, warm and cold fronts are most common in the Pacific and Atlantic cyclone tracks and the lee of the Rockies. In contrast with prior research, we find that the activity of warm and cold fronts is significantly modulated by the phase and intensity of El Niño–Southern Oscillation. The influence of El Niño is significant for winter warm fronts, winter cold fronts, and spring cold fronts, with activity decreasing over the continental United States and shifting northward with the Pacific and Atlantic cyclone tracks. Long-term trends are generally not significant, although we find a poleward shift in frontal activity during the winter and spring, consistent with prior research. We also identify a number of regional patterns, such as a significant long-term increase in warm fronts in the eastern tropical Pacific Ocean, which are characterized almost entirely by moisture gradients rather than temperature gradients.

1. Introduction

Frontal boundaries between air masses are an important trigger for precipitation and often serve as foci for severe thunderstorms (Sanders and Doswell 1995; Sanders 1999; Garner 2013; Catto et al. 2015; Dowdy and Catto 2017). The significance of fronts has been recognized since observation networks were first established over North America (e.g., Holzman 1937). Many climatologies have explored the regional and global occurrence of fronts and their links to significant weather, including their association with warm conveyor belts and parent extratropical cyclones (Serreze et al. 2001; Shafer and Steenburgh 2008; Payer et al. 2011; Simmonds et al. 2012; Catto and Pfahl 2013; Catto et al. 2015; Schemm et al. 2015; Dowdy and Catto 2017; Parfitt

et al. 2017; Schemm et al. 2018b; Thomas and Schultz 2019a). However, fronts have proven difficult to analyze objectively over long climatological periods. Early studies were highly sensitive to the subjective interpretation of meteorologists, especially with sparse surface observations, and required extensive human labor (Sanders and Doswell 1995; Sanders 1999; Sanders and Hoffman 2002).

Numerical frontal analysis (NFA) emerged in the 1960s (Renard and Clarke 1965; Clarke and Renard 1966) as a way to objectively define fronts. NFA algorithms apply human-conceived rules to analyze gridded data and determine which grid cells belong to a front. The most common procedure [summarized by Hewson (1998)] uses a locating variable, typically based on second or third derivatives of a thermal field, to identify all candidate fronts. Masking variables are then used to eliminate spurious fronts that do not meet user-specified

Corresponding author: John T. Allen, johnterrallen@gmail.com

DOI: 10.1175/JCLI-D-19-0680.1

© 2020 American Meteorological Society. For information regarding reuse of this content and general copyright information, consult the [AMS Copyright Policy](https://www.ametsoc.org/PUBSReuseLicenses) (www.ametsoc.org/PUBSReuseLicenses).

criteria (e.g., minimum length, translation speed, thermal gradient, or thermal advection). More recently, [Schemm et al. \(2018b\)](#) discussed a range of procedures and locating variables, arguing that the best practice is to use equivalent or wet-bulb potential temperature, although even using this field has limitations. Despite over 50 years of work in NFA, the algorithms are generally brittle—because of noise in derivatives computed on a finite grid—and have strong systematic biases. For example, [Schemm et al. \(2015\)](#) compared two popular methods and concluded that the thermal method rarely detects weakly baroclinic fronts (e.g., those induced by wind shear and convergence between two anticyclones), while the wind-shift method rarely detects warm fronts. In a similar vein, [Hope et al. \(2014\)](#) compared six NFA methods over southwestern Australia and identified substantial deficiencies in each method. Moreover, NFA is sometimes hampered by the nonuniform spacing of latitude–longitude grids (e.g., [Catto and Pfahl 2013](#)).

Recent work by [Lagerquist et al. \(2019\)](#), hereinafter [L19](#), demonstrates the advantages of using convolutional neural networks (CNNs) to identify fronts in gridded data. CNNs ([Fukushima 1980](#); [Fukushima and Miyake 1982](#)) combine traditional neural nets with specialized image-processing layers, called convolutional and pooling layers. A traditional neural net ([Goodfellow et al. 2016](#), chapter 6) contains only dense layers, which cannot represent spatial relationships. Thus, even when predictors are available on a grid, traditional neural nets must be trained with scalar features computed with summary statistics from the gridded data. Thus, they may not be optimal predictors for the given task. Conversely, in a CNN, features are created by the convolutional and pooling layers and then passed to the dense layers, which transform the features into predictions. The CNN simultaneously learns weights for creating the features and for transforming the features to predictions. This generally allows the CNN to pick more relevant features, which leads to better predictions. CNNs have been applied to a wide range of problems ([Dieleman et al. 2015](#); [Silver et al. 2016, 2017](#); [Rakhlin et al. 2018](#); [Karras et al. 2018](#); [Suwajanakorn et al. 2017](#)), but only recently have they become popular in the atmospheric sciences. Applications have included estimation of sea ice concentration from satellite images ([Wang et al. 2016](#)), extreme weather detection in climate models ([Racah et al. 2017](#); [Kurth et al. 2018](#)), detection of fronts ([L19](#)), replacement of parameterizations of subgrid-scale processes in climate models ([Rasp et al. 2018](#); [Brenowitz and Bretherton 2019](#); [Beucler et al. 2019](#)), and improving the prediction and understanding of various other weather phenomena ([Gagne et al. 2019](#); [McGovern et al. 2019](#)).

In addition to the advantages offered by CNNs, the ERA5 reanalysis ([Hersbach and Dee 2016](#)) has recently

become available. Unlike predecessors such as ERA-Interim, ERA5 is more appropriate for climatological studies because it uses consistent methods throughout the period (1979–2018) to estimate sea surface temperature, sea ice, greenhouse gases, and volcanic emissions. Recent evaluations of ERA5 ([Gossart et al. 2019](#); [Tarek et al. 2020](#)) suggest that ERA5 is similar in quality to other new reanalyses and appropriate for climatological studies. In this study we apply CNNs to 40 years (1979–2018) of ERA5 data to create a climatology of warm and cold fronts over North America. We describe mean characteristics for the entire period (1979–2018), trends over time, and variability with respect to the El Niño–Southern Oscillation (ENSO) state.

For ENSO, we test the hypothesis that El Niño and La Niña respectively cause an equatorward and poleward displacement in the mean position of fronts. El Niño and La Niña are characterized by anomalously warm and cold sea surface temperatures in the eastern equatorial Pacific, respectively, which result in the enhancement or suppression of deep convection. Previous work ([Seager et al. 2003](#); [Tandon et al. 2013](#); [Schmidt and Grise 2017](#)) suggests that enhancement of deep convection during El Niño causes strengthening and contraction of the Hadley cell, leading to an equatorward shift of the associated subtropical jet and midlatitude cyclone track. Conversely, suppression during La Niña causes the Hadley cell to weaken and expand, leading to a poleward shift of the jet and cyclone track. These effects are corroborated in a cyclone climatology ([Eichler and Higgins 2006](#)), studies of the jet position ([Allen et al. 2015](#); [Cook et al. 2017](#); [Schemm et al. 2018a](#)), and the frontal climatology of [Rudeva and Simmonds \(2015\)](#).

For long-term trends, we test whether expected responses to global warming are evident for midlatitude cyclones and their associated fronts. Arctic amplification ([Serreze and Barry 2011](#)) should lead to decreased low-level baroclinicity, and hence fewer cyclones and associated fronts at high latitudes. Second, some research suggests that global warming is driving an expansion of the Hadley cell toward the poles ([Davis and Rosenlof 2012](#); [Lucas et al. 2014](#); [Schmidt and Grise 2017](#)), which should lead to a poleward shift in the subtropical jet and cyclone track. Although the latitudes of the Hadley-cell edge and subtropical jet are highly correlated ([Kang and Polvani 2011](#)), there is disagreement on which is the driving factor. Traditionally it has been assumed that changes in the Hadley cell lead changes in the subtropical jet, but [Rudeva et al. \(2019\)](#) have suggested the opposite, with a lead time of 1–2 days. To our knowledge, only two studies have investigated long-term trends in cyclones or fronts over a large

TABLE 1. Predictor variables (all from ERA5 reanalysis).

Full name	Abbreviation	Vertical levels
Temperature	T	Surface, 1000, 950, 900, and 850 hPa
Specific humidity	q	Same
Zonal wind velocity	u	Same
Meridional wind velocity	v	Same
Wet-bulb potential temperature (AMS 2012)	θ_w	Same
Geopotential height	Z	1000, 950, 900, and 850 hPa
Pressure	p	Surface

(near-hemispheric to global) domain: Berry et al. (2011a) investigated fronts only, and Rudeva and Simmonds (2015) investigated both cyclones and fronts. These studies both found decreases in fronts over the Arctic, with no corresponding change to cyclone frequency, and northward shifts over the Pacific Ocean. In contrast, over the Atlantic Ocean there was an overall decrease, rather than a northward shift.

Confirming the above hypotheses would be a useful contribution to the literature, because 1) few previous studies have explored these links and 2) the previous studies used NFA methods that do not distinguish between warm and cold fronts. Section 2 describes the training and evaluation of CNNs; sections 3 and 4 describe methods and results, respectively, for climatological analysis; and section 5 summarizes and presents recommendations for future work.

2. Training and evaluation of CNNs

a. Preprocessing

Predictors come from the ERA5 reanalysis, and labels, treated as correct answers, come from Weather Prediction Center (WPC) surface bulletins (<https://www.wpc.ncep.noaa.gov/html/sfc2.shtml>). The bulletins include polylines identifying warm and cold fronts every 3 h (0000, 0300, . . . , 2100 UTC daily). We have obtained bulletins for ~ 9.2 years (Table 2). An alternative approach would be to use NFA-derived labels instead, because unlike human labels, they are consistent and based directly on a priori scientific reasoning. However, humans outperform NFA at pattern recognition, and since the WPC labels come from an ensemble of trained meteorologists, they are ultimately based on scientific reasoning as well. Also, NFA often has strong systematic biases [see discussion of Schemm et al. (2015) in section 1], which could be easily overfit by a CNN. The WPC labels do not appear to contain such biases, likely because each bulletin is a consensus among many meteorologists and the set of meteorologists for each

bulletin is different. Thus, we believe that training a CNN with WPC labels is the best possible approach, because it leverages the pattern-recognition ability and scientific knowledge of humans.

The ERA5 reanalysis is available at 1-h time steps on a global 0.281° grid. For training and evaluation of the CNN, we use only 3-hourly data (coinciding with WPC bulletins), downloaded from the Climate Data Store (<https://doi.org/10.24381/cds.bd0915c6>). We use the variables listed in Table 1, which hereinafter will be abbreviated as in Table 1.

Before training, we process the ERA5 and WPC data as follows.

- 1) Interpolate ERA5 data from the native latitude–longitude grid to an equidistant grid. This is necessary because CNNs implicitly assume equidistant grids. We use the 32-km NARR grid, as in L19, but extending 100 grid cells farther in each direction, yielding dimensions of 477×549 .
- 2) Rotate winds from Earth-relative to grid-relative. This puts wind vectors and scalar fields in the same coordinates, which makes relevant quantities such as temperature and moisture advection easier for the CNN to represent. This decision was based on a priori reasoning. As we did not also attempt training CNNs with Earth-relative winds, it is also possible that CNNs trained with Earth-relative winds would perform just as well.
- 3) Normalize predictors. Specifically, at each time step, we apply the following equation to each variable at each vertical level:

$$z = \frac{x - \bar{x}}{s}, \quad (1)$$

where x is the original value, $\bar{x}$ is the mean over the full grid, s is the standard deviation over the full grid, and z is the normalized value. Normalized values follow a Gaussian distribution, with mean of 0.0 and variance of 1.0. This prevents predictors with large values from causing large

FIG. 1. Masking of grid cells for CNN training and evaluation: (a) Number of warm fronts in the WPC dataset (1500 UTC 5 Nov 2008–2100 UTC 31 Dec 2017) after dilation. (b) Number of cold fronts after dilation. (c) The resulting mask, based on which grid cells have at least 100 total fronts [sum of (a) and (b)]. Grid cells outside the green area are masked and cannot be used as the center of a training, validation, or testing patch.

weight updates in the CNN, which could cause weights to oscillate rather than converge to optimal values.

- 4) Convert WPC fronts from polylines to the 32-km grid. Each grid cell is labeled WF if intersected by a warm front, CF if intersected by a cold front, and NF otherwise.
- 5) Dilate WPC fronts, effectively replacing each frontal grid cell with a 3×3 neighborhood (see L19 for details).
- 6) Mask grid cells where the WPC does not label fronts. Masked grid cells (white in Fig. 1c) cannot be used as the center of a training, validation, or testing patch, because their true labels are unknown.
- 7) Create patches (one for each unmasked grid cell at each time step). Each patch contains the predictors on a 33×33 grid ($1056 \text{ km} \times 1056 \text{ km}$) and postdilation label (NF, WF, or CF) at the center grid cell.
- 8) Split patches into training, validation, and testing sets (Table 2).

b. Description of CNNs

This section briefly describes the inner workings of CNNs. See L19 for a more thorough description.

The main components of a CNN are convolutional and pooling layers, which transform the inputs into abstractions called “feature maps,” and dense layers, which transform

the feature maps into predictions. Figure 2 shows the architecture of a CNN used in this study.

A convolutional layer uses convolutional filters to transform the input maps. Convolutional filters have been used in image-processing for decades to perform operations such as blurring, sharpening, and edge detection (Shapiro and Stockman 2000, chapter 5). In these applications the filter weights are fixed, but in a CNN they are learned over time. Convolutional layers generally increase the number of channels, which increases the number of abstractions that can be subsequently leveraged for prediction. For example, in Fig. 2, the input maps have eight channels (all weather variables) and feature maps produced by the first convolutional layer

TABLE 2. Training, validation, and testing periods for CNN. There is a 1-week gap between each pair of consecutive datasets to ensure that they are independent (not temporally autocorrelated).

Dataset	Time period
Training	1500 UTC 5 Nov 2008 –2100 UTC 24 Dec 2014
Validation	0000 UTC 1 Jan 2015 –2100 UTC 24 Dec 2016
Testing	0000 UTC 1 Jan 2017 –2100 UTC 31 Dec 2017

FIG. 2. Example of CNN architecture. The input is a 33×33 grid with eight predictors (T , q_w , u , and v at the surface and 850 hPa). (a), (b) Predictors: heat maps are θ_w , and gray lines are wind barbs, from the ERA5 reanalysis. Dark blue circles are warm fronts, and light blue triangles are cold fronts, from the WPC bulletins. (c) Feature maps produced by the first convolutional layer, after activation and batch normalization. Red values are positive; blue values are negative. (d) Feature maps produced by the first pooling layer. (e) Feature maps produced by the third convolutional layer. (f) Feature maps produced by the second pooling layer. The flattening layer transforms maps from the last pooling layer into a vector of length 4608 ($8 \times 8 \times 72$), and the dense layers transform this vector into representations of exponentially decreasing length ($4608 \rightarrow 399 \rightarrow 35 \rightarrow 3$), terminating with the prediction vector. The prediction vector contains probabilities of the three classes: no front, warm front, and cold front.

have 36 channels, each a transformation of the eight original channels. Convolution is a spatial and multivariate operation, so it is performed over all spatial dimensions and all channels at the same time. Thus, each channel after the input layer contains information from all eight original weather variables.

Each convolutional layer includes two operations after the convolution itself: activation and batch normalization (Ioffe and Szegedy 2015). Activation is a nonlinear function applied elementwise to the feature maps. Without this nonlinearity, the CNN could learn only linear relationships, because convolution is a linear operation and any series of linear operations is linear. We use the leaky rectified linear unit (ReLU; Maas et al. 2013) with a slope parameter of 0.2. Activation is followed by batch normalization, which transforms each element (each channel at each grid point) to approximately a standard Gaussian distribution, with mean of 0.0 and variance of 1.0. This

helps the CNN to learn more quickly and mitigates the vanishing-gradient problem (discussed in L19).

A pooling layer coarsens feature maps, using either a maximum or mean filter. Following common practice, this work uses a maximum filter with coarsening factor of 2, which halves the spatial resolution (doubles the grid spacing). Thus, pooling layers in Fig. 2 increase the grid spacing from 32 to 64 to 128 km, which allows deeper convolutional layers to learn larger-scale features. This is advantageous because weather analysis and prediction often depend on features at multiple scales.

The dense layers are spatially agnostic, so feature maps are flattened into a 1D vector before they are passed to the dense layers. Dense layers transform this vector into representations of decreasing length (exponentially decreasing in Fig. 2), terminating with the prediction vector. This transformation is a nonspatial version of the linear transformation performed by

convolutional layers. All dense layers except the last follow this transformation with a leaky ReLU activation (slope of 0.2) and batch normalization, like the convolutional layers. The last dense layer uses the softmax activation function (section 6.2.2.3 of Goodfellow et al. 2016), which forces all outputs to range over $[0, 1]$ and forces their sum to 1.0, allowing them to be interpreted as probabilities of mutually exclusive and collectively exhaustive classes.

Adjustable weights reside in the convolutional and dense layers. These weights are initialized randomly and then adjusted during training to minimize the cross-entropy between predicted and true labels. To reduce overfitting, we use L_2 regularization (section 4.4.2 of Chollet 2018) with a strength of 0.001 for the convolutional layers and dropout regularization (Hinton et al. 2012) with a rate of 50% for all dense layers except the last. Dropout, used during the training phase only, randomly omits the specified fraction of weights for each example. This causes redundancy of weights in the same layer, which reduces overfitting. However, when applied to the last dense layer (the output layer), dropout often changes the predictions dramatically, leading to a decrease in skill.

c. Finding the best CNN

We conduct a grid search to find the best hyperparameters for the CNN. A “hyperparameter” is a value that, unlike weights in the convolutional and dense layers, cannot be adjusted during training and therefore must be chosen a priori. “Grid search” (section 11.4.3 of Goodfellow et al. 2016) means that we try all 288 combinations of the values listed in Table 3. A key difference from L19 is that we train some CNNs with data from two vertical levels. We hypothesized that this would improve skill, since fronts are 3D phenomena despite the common practice of treating them as 1D or 2D. Every CNN includes 1000-hPa or surface data, because L19 find that 1000 hPa is the best single level and WPC fronts are nominally based on surface data. A “convolutional block” is a series of one or more convolutional layers, followed by a pooling layer (e.g., the CNN in Fig. 2 contains two blocks with two layers each). The number of blocks and layers are important because they control computing time for applying the CNN. For the CNNs trained in this experiment, computing time ranges from 10 to 30 min for the full grid at one time step. There are 116 880 three-hourly time steps in the climatology, so computing time for the climatology ranges from 19 480 to 58 440 core-hours. Even when parallelized over 90 cores, the resulting clock time ranges from 9 to 27 days. Thus, in choosing the “best” CNN, we must consider size as well as predictive skill. Long computing times

TABLE 3. Hyperparameters for CNN. Each set of predictor variables is formed by including or excluding the fundamental thermal fields (T and q), including/excluding θ_w , and including/excluding pressure information (p for the surface level; Z for isobaric levels).

Hyperparameter	Values attempted
Predictor variables	u , v , and θ_w u , v , T , and q u , v , θ_w , T , and q u , v , θ_w , and p or Z u , v , T , q , and p or Z u , v , θ_w , T , q , and p or Z
Vertical levels	Surface 1000 hPa Surface and 1000 hPa Surface and 950 hPa Surface and 900 hPa Surface and 850 hPa
No. of convolutional blocks	2 and 3
No. of convolutional layers per block	1, 2, 3, and 4

also explain why we use 3-hourly, rather than hourly, time steps to create the climatology.

As in L19, each CNN is trained with 3200 batches and batches are downsampled so that each contains 50% NF patches, 25% WF, and 25% CF. Downsampling is not done for the validation or testing data, so these results are based on the real-world distribution (98.95% NF, 0.27% WF, and 0.78% CF, after dilation). Also as in L19, CNNs are ranked by their Gerrity score on 1 million patches selected randomly from 1000 time steps in the validation period (Fig. 3). The Gerrity score ranges over $[-1, 1]$, and higher values are better. There is no coherent trend with respect to predictor variables or architecture, but there is a trend with respect to vertical levels. Models trained with surface data plus isobaric data from 950 hPa or aloft perform better than those trained with surface data only, 1000-hPa data only, or both. Because of computing time, the experiment is limited to the vertical levels shown in Table 3, based on prior experience and recommendations from the literature (Schemm et al. 2018b; Thomas and Schultz 2019a). Peirce and Heidke scores (defined in L19 but not shown here) generally follow the same pattern as Gerrity scores in Fig. 3.

We choose the CNN with the second-best Gerrity score (0.6827) because it is smaller and more computationally efficient than that with the best score (0.6835). The chosen CNN has two blocks, two layers per block, and is trained with surface and 850-hPa fields of T , q , u , and v . Ranking by Gerrity score and downsampling the training data both encourage overprediction of fronts, but this overprediction is largely mitigated by conversion to frontal zones (section 2d).

FIG. 3. Gerrity score on the validation period for all combinations of CNN hyper-parameters. The CNN with the best and second-best Gerrity scores are identified with a black diamond and black circle, respectively.

FIG. 4. Results of permutation test for best CNN. For the single-pass test, the bar for variable x shows model performance when only x is permuted (all other variables are clean or intact). For the multipass test, the bar for variable x shows model performance when variable x and all variables above it are permuted. In both cases, variable importance decreases from top to bottom (the most important is at the top).

We use the permutation test (McGovern et al. 2019) to rank predictor importance for the best CNN. The importance of predictor x is determined by how much model performance declines when x is permuted (i.e., when maps of x are randomly shuffled over all patches). We run the test on 50 000 patches, sampled randomly from 1000 time steps in the testing period, using multiclass area under the receiver-operating-characteristic curve (AUC; Metz 1978) as a fitness function. In the single-pass test (Fig. 4), in which only one predictor is permuted at once, the five most important predictors are wind components and 850-hPa temperature, while the two least important are specific humidities. In the multipass test, where

multiple predictors are permuted at the same time, once surface wind components have been permuted, 850-hPa wind components appear less important because they are highly correlated with their surface analogs (Pearson correlation of 0.77 between surface and 850-hPa u wind; Pearson correlation of 0.78 between surface and 850-hPa v wind). This is why thermal variables appear to be more important in the multipass test.

d. Postprocessing (conversion to frontal zones)

The CNN produces three probabilities at each grid cell (Fig. 5b). For the purpose of the climatology, we convert these to deterministic objects [frontal zones

FIG. 5. Determinization procedure: (a) Surface predictors at 1800 UTC 2 Oct 2016. Formatting is explained in Fig. 2. (b) CNN-estimated probabilities of WF (red shading) and CF (blue shading). (c) Deterministic predictions, using thresholds $p_{WF}^* = p_{CF}^* = 0.65$ in Eq. (2). (d) As in (c), but excluding small regions (major axis < 200 km) that are not within 200 km of a large region.

similar to those in Schemm et al. (2018b)], using the following procedure at each time step.

- 1) Convert probabilities to deterministic labels, using Eq. (2). The p_{WF} and p_{CF} are the CNN-estimated WF and CF probabilities, respectively; p_{WF}^* and p_{CF}^* are thresholds on these probabilities, and $\hat{y}$ is the resulting deterministic label:

$$\hat{y} = \begin{cases} \text{CF} & \text{if } p_{CF} \geq p_{CF}^* \text{ and } p_{CF} \geq p_{WF} \\ \text{WF} & \text{if } p_{WF} \geq p_{WF}^* \text{ and } p_{WF} > p_{CF} \\ \text{NF} & \text{otherwise} \end{cases} \quad (2)$$

Subjectively, we have found that $p_{WF}^* = p_{CF}^* = 0.65$ yields realistic frontal zones (Fig. 5c). Lower thresholds yield frontal zones that are too wide, while higher thresholds yield frontal zones that are too narrow and often miss fronts altogether. However, we have found that thresholds of 0.50 and 0.80 yield qualitatively the same climatology (spatial patterns and subsequent conclusions vis-à-vis ENSO and long-term climate change), although with higher frontal frequencies for the 0.50 threshold and lower frequencies for the 0.80 threshold.

- 2) Find frontal regions (sets of connected grid cells with the same type, either WF or CF). For example, in Fig. 5c three connected regions (all CF) touch the continental United States (CONUS).
- 3) Throw out small WF regions (major axis < 200 km) that are not within 200 km of a large WF region. Do the same for CF regions. When a small region is close to a large region, they are often part of the same frontal zone but separated by a “hole” with weaker frontal properties such as thermal gradient or advection. We use a 200-km threshold, because it is the minimum length for the synoptic scale (Orlanski 1975) and the climatology is geared toward synoptic-scale fronts. We have found that thresholds of 400 and 600 km make very little difference to the climatology or subsequent conclusions, either qualitatively or quantitatively, because the vast majority of CNN fronts are longer than 600 km. In Fig. 5d the small CF region off the California coast is eliminated, while the small CF region in Saskatchewan is kept, because the latter is within 200 km of a large CF region.

We use four scores (defined in the appendix) to evaluate frontal zones for the testing period: POD, FAR, CSI, and frequency bias. These are computed

FIG. 6. Performance diagram (Roebber 2009) on testing data for the selected CNN with the selected probability thresholds [$p_{WF}^* = p_{CF}^* = 0.65$ in Eq. (2)]. Dashed gray lines show frequency bias. All scores (POD, FAR, CSI, and frequency bias) are defined in the appendix. Each point represents one neighborhood distance. Error bars show the 95% confidence interval, determined by bootstrapping (Efron 1979) the testing data 1000 times. The error bars are nearly invisible because the confidence intervals are very narrow; e.g., at 150 km the interval for CSI is [0.3720, 0.3727].

with a neighborhood distance, used to match predicted (CNN) with actual (WPC) fronts. Results at four neighborhood distances are shown in Fig. 6. The frequency bias is positive at all neighborhood distances, suggesting that the CNN detects too many fronts. However, in our judgement the WPC often misses fronts (e.g., in Fig. 5, the cold front south of Nova Scotia and pair of fronts in the Canadian Arctic). Thus, we consider a positive frequency bias acceptable.

3. Methods for climatological analysis

We apply the chosen CNN to the full 477×549 grid, at all 3-hourly time steps from 1979 to 2018. Although some grid cells were masked out for training and evaluation, no grid cells are masked in this phase, because the CNN has already been trained and WPC labels are no longer needed to provide the CNN with correct answers. We segregate each analysis by season: winter [December–February (DJF)], spring [March–May (MAM)], summer [June–August (JJA)], and autumn [September–November (SON)]. We analyze four properties: WF frequency, CF frequency, WF length, and CF length. WF (CF) frequency is defined as the fraction of time steps with a warm (cold) front, and length is defined as the major-axis length of the frontal zone. Length is computed only for fronts

that do not touch the edge of the grid, because in these cases there is no way to ascertain how far the front extends beyond the edge.

We also explore variability and long-term trends in frequency and length. For variability, we define ENSO state by the standardized anomaly of the Niño-3.4 index, which is computed with sea surface temperature anomalies relative to 1981–2010 (Climate Prediction Center 2019). Each comparison—between a composite of ENSO-neutral months and a composite of nonneutral months—is done for one season, so that ENSO signals are not confounded with seasonal signals. The nonneutral composite contains months with El Niño, La Niña, strong El Niño, or strong La Niña conditions, all defined in Table 4. Statistical significance is determined by a two-tailed Monte Carlo test with 20 000 iterations.

To compute long-term trends and determine their significance, we use a Mann-Kendall test (Mann 1945; Kendall 1955), also segregated by season. We use the pyMannKendall package (Hussain and Mamhud 2019), which returns two values: the slope of the Theil-Sen line (Sen 1968; Theil 1992) and the two-tailed p value.

When conducting multiple hypothesis tests (in this case, one at each grid point), one must control the false discovery rate (FDR). For example, if each test were conducted independently with a 95% confidence level, one would expect 5% of grid points to show “significant” differences even with purely random data. For both ENSO-related variability and long-term trends, we use Eq. (3) from Wilks (2016), which decreases the p value threshold (p^*) to ensure an upper bound on the FDR. Specifically, we set the maximum FDR to 0.10, which typically yields a p^* on the order of 10^{-3} or lower. Only grid points with $p \leq p^*$ are considered significant.

4. Results and discussion

a. Climatology of warm and cold fronts

Winter cold fronts (Fig. 7a) over the central and eastern Pacific are predominantly confined to a band from 25° to 50°N , with the highest number from 30° to 40°N . This swath shifts northward as it approaches the west coast of North America, along the northeast Pacific cyclone track (Eichler and Higgins 2006; Hodges et al. 2011). Most of these fronts dissipate as they encounter high terrain or southern California. Landfalling cold fronts on the west coast are most common in the Pacific Northwest, and fronts reaching the interior of western North America are most common in the desert southwest, where the terrain is less steep. Over the continent, cold fronts are most common from the lee of the Rockies to Appalachia, consistent with this belt of cyclonic activity (Hodges et al. 2011; Schemm et al. 2018a). The

TABLE 4. ENSO-based composites for Monte Carlo test, by season. Each table cell contains the number of months (from 1979 to 2018) in the given composite; z is the standardized anomaly of the Niño-3.4 index.

Composite	Winter (DJF)	Spring (MAM)	Summer (JJA)	Autumn (SON)
Neutral ($ z < 0.5$)	38	71	71	50
El Niño ($z \geq 0.5$)	35	23	24	34
La Niña ($z \leq -0.5$)	47	26	25	36
Strong El Niño ($z \geq 1.0$)	22	11	9	20
Strong La Niña ($z \leq -1.0$)	18	6	8	21

southern extent of regular CF occurrence is the Yucatán Peninsula of Mexico. Over the Atlantic, CF are generally east of the Gulf Stream, following the east-coast cyclone track (Eichler and Higgins 2006; Schemm et al. 2018a). There are also regional maxima of ≥ 0.08 over New Mexico and Texas; over eastern Appalachia and the east coast, where cold-air damming is frequent (Rackley and Knox 2016); and in the northeast Gulf of Mexico, associated with extension of the Gulf Coast cyclone track (Schemm et al. 2018a) during winter.

Winter WF (Fig. 8a) are less frequent than CF over the continent, reflecting the seasonal availability of warm advection (Schemm et al. 2018b). As would be expected, the most prominent WF maxima occur north of corresponding CF maxima, consistent with mean frontal positions relative to their parent cyclones. Warm-frontal maxima are found in the central Pacific,

off the west coast from Washington State northward, in the lee of the Rockies, and in the Atlantic southeast of the Canadian Maritimes and Newfoundland. There are also maxima near the Gulf Coast and Great Lakes, associated with warm advection from moist tropical air masses (Henry 1979; Crisp and Lewis 1992; Payer et al. 2011).

Spring and autumn CF (Figs. 7b,d) have similar distributions, but autumn maxima are less prominent and displaced slightly northward. Relative to winter, spring and autumn in both ocean basins are shifted northward by $\sim 10^\circ$, consistent with the seasonal migration of the cyclone track (Hodges et al. 2011). This shift has a strong influence on frontal activity over the continent, leading to fewer fronts in the Gulf of Mexico and more east of the Rockies, the latter associated with stronger low-level baroclinicity over the continent. In spring this

FIG. 7. Frequency of cold fronts in each season from 1979 to 2018. Frequency is the fraction of 3-h time steps with a cold front.

FIG. 8. As in Fig. 7, but for warm fronts.

continental maximum exceeds those in the ocean basins, whereas in the autumn continental and oceanic maxima are roughly equal in magnitude. There are also regional maxima in the lee of the Appalachians, over southeast Nevada, and from central to eastern Canada, with the highest values southwest of James Bay. This last maximum may be associated with the northern part of the Alberta Clipper cyclone track (Thomas and Martin 2007).

Spring WF (Fig. 8b) generally occur north of spring CF, and far fewer WF occur over the continent. However, there are two exceptions: a diffuse maximum from the Plains to Lake Ontario and a sharper maximum south of Hudson Bay, consistent with earlier analyses from subjectively analyzed fronts (Morgan et al. 1975). The first reflects persistent WFs that lift northward during periods of warm advection and then retreat southward overnight (Payer et al. 2011), while the second reflects warm air moving over landfast sea ice. Meanwhile, the main difference in autumn warm-frontal activity (Fig. 8d), relative to cold fronts, is a northward shift and sharp maximum along the west coast, the latter associated with seasonally warm water relative to the more rapidly cooling land.

During summer, CF (Fig. 7c) over the ocean are less common than in other seasons and generally occur north of 35°N. CF over the continent are generally confined to similar latitudes but occur more often. In Canada and the northern Plains, CF activity peaks during the

summer, reflecting the annual heating cycle and the poleward shift of low-level baroclinicity. Curiously, there is also a maximum over inland parts of the Pacific Northwest. In contrast, summer WF (Fig. 8c) over the continent are generally confined to the northern Plains and north coast. The maximum along the north coast, associated with seasonally warm land relative to the more slowly warming ocean, disappears in the autumn (Fig. 8d), when the ocean reaches its annual temperature peak and the land again starts to cool. Also, warm fronts have a summer and autumn maximum in the tropical eastern Pacific. Inspection of individual cases reveals that these fronts are characterized almost entirely by moisture gradients rather than temperature gradients, due to a boundary between moist tropical air and dry subtropical air. Berry et al. (2011a) found a similar maximum in the same area and drew similar conclusions. In addition, some of these warm fronts are associated with tropical cyclones undergoing extratropical transition (Camargo et al. 2008; Wood and Ritchie 2014).

We also explore the mean length of WF and CF (Figs. 9 and 10). CF are generally twice as long as WF, reflecting the stronger thermal gradients associated with the former. During winter, four maxima in warm-frontal length stand out (Fig. 10a): the central to eastern Pacific, east of the Gulf Stream, along the east coast of the United States, and in the lee of the Rockies. The last two

FIG. 9. Average cold-frontal length (km) in each season from 1979 to 2018. The length of each front is computed separately at each time step, and the value at grid cell g is the average for all cold fronts touching g . For each season, grid cells with <100 cold fronts are not plotted, so as to reduce clutter.

are associated with often explosive development of extratropical cyclones (Allen et al. 2010; Hodges et al. 2011; Schemm et al. 2018a) and lee cyclogenesis, respectively. In the last three areas, peak warm-frontal activity occurs in the winter; in the central to eastern Pacific, peak activity occurs in the spring and summer, when northward transport of warm air is stronger. The climatology of cold-frontal length is very different. During the winter (Fig. 9a), the highest values (mean lengths up to 3000 km) occur along the east coast of the United States and in the Atlantic and Caribbean. These are associated with long trailing cold fronts separating cold continental air from warm tropical air, which often stretch deep into the tropics (Henry 1979). CF over the Pacific and the continent are generally much shorter. Patterns in the other seasons (Figs. 9b–d) are very similar to winter, except that cold fronts are shorter, due to the annual cycle and diabatic heating depriving the Northern Hemisphere of cold air.

b. Variability

To explore the role of climate variability in modulating frontal frequency, we consider the influence of ENSO. A winter El Niño of any strength (weak or strong) creates a large response in both warm- and cold-frontal activity (Figs. 11a,b). The WF frequency

increases significantly over parts of the Pacific from 30° to 40° N, while decreasing from 45° to 60° N, reflecting a southward shift in activity. Typical changes in WF frequency are ± 0.012 , whereas frequencies averaged over the 40 years are ~ 0.08 , representing a relative change of $\sim 15\%$. This is consistent with our hypothesized southward shift during El Niño winters, also identified in Eichler and Higgins (2006) and Rudeva and Simmonds (2015). Over the continent, southward displacement of cold air reduces continental low-level baroclinicity, causing the cyclone track in the lee of the Rockies to shift southward (Eichler and Higgins 2006; Allen et al. 2015; Schemm et al. 2018a). This leads to significant decreases in WF frequency in the western Plains. The main difference between the WF and CF responses is that the CF response is more spatially coherent and significant. Typical changes in CF frequency are 0.015–0.020, while typical averages over the period are 0.10, representing a relative change of 15%–20%. Over the continent, CF frequency decreases significantly in much of the CONUS, unlike WF frequency, which decreases significantly only in the western Great Plains. This reflects cold advection in response to the subtropical-jet-driven activation of the Gulf Coast cyclone track, which causes CF to push farther southward, leading to the increase seen in the Gulf of Mexico (Eichler and Higgins 2006; Allen et al. 2015; Schemm et al. 2018a).

FIG. 10. As in Fig. 9, but for warm fronts.

To determine if these signals are modulated by the strength of El Niño, we also consider strong El Niño only (Figs. 12a,b). For both WF and CF, the pattern is similar to that caused by any El Niño (cf. Figs. 11a,b and 12a,b), except that differences are larger and more broadly significant. Increasing the threshold to strong El Niño causes a greater change for WF than for CF, which may indicate stronger advection of air masses and cyclonic activity with increasing El Niño strength (Schemm et al. 2018a). Overall, in contrast with Hardy and Henderson (2003), we conclude that WF and CF activity, over both North America and the surrounding oceans, are strongly modulated by winter El Niño.

For spring El Niño of any strength (Figs. 11c,d), results are weaker than for winter. Significant changes are found only for CF and predominantly over the eastern Pacific, where the pattern represents a southward shift as in winter. The weaker signal in spring is not surprising: ENSO's period of peak influence is the autumn and winter, typically followed by a gradual transition to neutral conditions during the spring and a weakening of atmospheric teleconnections (Lee et al. 2014). However, during strong spring El Niño the signal for CF is more robust, with significant increases in the subtropical Pacific, including the Gulf of California, and in the Caribbean. CF frequency also increases over the southern Great Plains and southeast CONUS,

consistent with the hypotheses of Allen et al. (2015) of a persistent offshore flow regime. CF decreases farther north are again related to a lack of cyclones in the lee of the Rockies (Schemm et al. 2018a), which typically support the development of frontal zones in this area. Instead, CF are confined farther south, with the baroclinic zone that promotes extratropical cyclogenesis along the Gulf Coast (Eichler and Higgins 2006; Schemm et al. 2018a).

Summer and autumn responses to El Niño (Figs. 11e–h and 12e,f) are generally not significant, caused by smaller sample sizes (Table 4) and weaker atmospheric teleconnections, especially during summer. However, there are two noteworthy signals. Autumn El Niño (Figs. 11g and 12g) causes an increase in WF frequency along the Gulf Coast and off the Atlantic coast of the southeastern CONUS, while strong autumn El Niño (Fig. 12g) causes a decrease in WF frequency in the western Great Plains. Both of these signals are similar to those found in winter (Figs. 11a and 12a).

For a winter La Niña of any strength (Figs. 13a,b), significant changes occur almost exclusively for WF. WF frequency increases significantly over the Pacific, Bering Sea, and midwestern United States. Increases over the Pacific are associated with a poleward displacement of ridging over the west coast, which drives the extratropical cyclone track and thus warm fronts northward (Eichler

FIG. 11. The WF and CF frequency (defined as in Figs. 7 and 8) during El Niño. The value shown in each panel is the composite difference (mean over El Niño months minus mean over ENSO-neutral months). Stippling shows where this difference is significant at the 95% confidence level, determined by a two-tailed Monte Carlo test with 20 000 iterations. In each panel, grid cells with <100 warm fronts (between the two composites: El Niño and neutral) are not plotted, so as to reduce clutter.

and Higgins 2006). In the midwestern United States, strengthening of the subpolar jet leads to downstream troughing over the Rockies, which enhances cyclogenesis in the Rocky Mountain cyclone track (Schemm et al.

2018a). However, strong La Niña (Fig. 14a) does not enhance these signals (i.e., leads to very few significant grid points), which may be caused by small sample size (Table 4).

FIG. 12. As in Fig. 11, but only for strong El Niño.

For a spring La Niña of any strength (Figs. 13c,d), significant changes occur almost exclusively for CF, which decrease over the central Pacific and increase offshore of the Pacific Northwest. Similar to spring El Niño (Fig. 11d), La Niña also promotes marginally robust CF increases over the southern CONUS, suggesting an increase in cyclonic activity east of the Rockies (Schemm et al. 2018a). Like El Niño, summer and autumn La Niña patterns (Figs. 13e–h and 14e,f) are generally not significant, caused by smaller sample size and weaker teleconnections.

c. Long-term trends

Some previous studies have shown modest or negligible trends in frontal frequency over North America in the last ~40 years (Berry et al. 2011a; Rudeva and Simmonds 2015). For example, Rudeva and Simmonds

(2015) found a significant decreasing trend in winter frontal frequency over Canada, an increase over the southern CONUS, and a poleward shift over the Pacific cyclone track. Our results are similar (Figs. 15a,b), although unlike Rudeva and Simmonds (2015) we do not find these changes to be significant outside the CONUS. This may be due to the fact that we consider WF and CF separately, while Rudeva and Simmonds (2015) consider both frontal types together, increasing their sample size for significance testing. We find poleward shifts in frontal activity over both the Pacific and Atlantic cyclone tracks, consistent with the findings of Berry et al. (2011a) on several reanalysis datasets, but in our case this result is not statistically significant. We also find significant increases in frontal activity (especially for CF) over the CONUS, which suggests an increase in cyclonic activity. However, the drivers of this signature are

FIG. 13. As in Fig. 11, but for La Niña.

unclear, as previous studies (e.g., [Rudeva and Simmonds 2015](#)) have not found a corresponding increase in cyclones. The increase we find could be due to weak cyclones, which are not well detected over complex terrain ([Neu et al. 2013](#)); cyclones moving more slowly, so that fronts reside over the continent for longer period; or perhaps a small northward migration of the Gulf Coast cyclone track. Increases in both warm- and cold-frontal

activity are strongest and most robust over the North Atlantic, where there is also some evidence that cyclone frequency has increased concurrently ([Rudeva and Simmonds 2015](#)).

Outside winter, there are few significant trends in frontal frequency (Figs. 15c–h). The CF trend in spring (Fig. 15d) is similar to the CF trend in winter, though less broadly significant, and the WF trend in summer

FIG. 14. As in Fig. 13, but only for strong La Niña.

(Fig. 15e) includes a significant increase in the eastern tropical Pacific, associated primarily with gradients of moisture rather than temperature. We note that the field-based significance test (discussed in section 3) tends to penalize small areas of change (Wilks 2016), which may mask other significant changes in Figs. 11–16.

We also consider long-term trends in frontal length. In the winter and spring (Figs. 16a–d), significant changes occur mainly in the subtropics, where the sample size is small, and the sign of the change (positive or negative) is not spatially coherent. Outside of these changes, winter CF length (Fig. 16b) has increased in the western CONUS while summer CF length (Fig. 16f) has decreased over the southeastern CONUS. Combining this information with frequency trends (Fig. 15), we suggest that the decrease in summer is driven by a decrease in cold air reaching the southeastern CONUS, while the increase in winter is

driven by more cold air reaching the western CONUS, allowing cold fronts to be well defined farther south.

5. Summary and future work

We used a convolutional neural network to develop an objective 40-yr climatology of warm and cold fronts over North America. The CNN was trained on ERA5 reanalysis fields to emulate fronts drawn by human meteorologists. The optimal predictors were surface and 850-hPa fields of temperature, specific humidity, u wind, and v wind. Combining surface and 850-hPa data is consistent with the detection-layer recommendations of Hewson (1998) and Sanders and Hoffman (2002), as well as discussions in Schemm et al. (2018b). We applied the trained CNN to 40 years of 3-hourly ERA5 analyses over the North American continent and adjacent ocean

FIG. 15. Linear trend in frontal frequency (per 40 years) in each season from 1979 to 2018. Frequency is the fraction of 3-h time steps with a front of the given type (warm or cold), as in Figs. 7 and 8. In each panel, the trend is determined by a Theil–Sen fit. Stippling shows where the trend is significant at the 95% confidence level, according to a two-tailed Mann–Kendall test. In each panel, grid cells with <100 fronts of the given type are masked out.

basins. Like Schemm and Sprenger (2015) and Schemm et al. (2018b), we represented fronts as 2D regions, which is a more realistic characterization than the 1D lines that are often used. From this climatology, we derived three analyses of the resulting fronts.

The first, which considers long-term averages of frontal frequency and length over the 40 years, differs from previous studies by considering warm and cold fronts separately. Our climatology identifies regional patterns not found in other studies (e.g., Berry et al. 2011b; Rudeva and

FIG. 16. Linear trend in frontal length (kilometers per 40 years) in each season from 1979 to 2018. As in Fig. 15, the trend is determined by a Theil–Sen fit; stippling shows significance at the 95% level, and grid cells with <100 fronts are masked out.

Simmonds 2015; Thomas and Schultz 2019a,b), such as cold fronts over the Intermountain West (Shafer and Steenburgh 2008) and a spring warm-frontal maximum over the lower Great Lakes (Payer et al. 2011). Second, we examined ENSO influences on warm- and cold-frontal frequency, finding that the signals are strongest and most robust in the winter and spring. Our results contrast with

Hardy and Henderson (2003) and Rudeva and Simmonds (2015), who suggested that ENSO has minimal influence on frontal frequency over North America. However, results are somewhat consistent with storm-track analyses by Eichler and Higgins (2006) and Schemm et al. (2018a), the latter having found that ENSO modulates the extratropical-cyclone track in winter. Unlike these two

studies, we found significant ENSO modulation of frontal activity in the spring, consistent with the synoptic analyses of Allen et al. (2015). Third, we examined long-term trends, finding little statistical significance. However, we found a general poleward shift in frontal activity during the winter and spring, consistent with Rudeva and Simmonds (2015) and Berry et al. (2011a). We also found a significant increase in winter cold fronts over the CONUS, which may be due to fronts moving more slowly over the continent. Further evidence, tied to cyclone frequency and other processes in ERA5, is needed to further investigate the drivers of this change.

As noted in section 2d, we have explored the sensitivity of climatological results to some hyperparameters of the detection method. Specifically, we obtained nearly identical results when increasing the minimum frontal length from 200 to 400 or 600 km, and results are qualitatively similar when making the probability threshold 0.50 or 0.80 instead of 0.65. Also, we explored running the CNN on half-resolution data (64-km spacing) and quarter-resolution data (128-km spacing). Although we did not have enough computing resources to redo all climatological analyses with the coarser data (this would require rerunning the CNN twice on the 40-yr period), we computed warm- and cold-frontal frequency for eight years (1980, 1985, . . . , 2015) with the full-resolution, half-resolution, and quarter-resolution data. The results were nearly identical, which suggests that our analyses do not depend on high spatial resolution. One area of particular sensitivity is for CF over the Intermountain West (Fig. 1), where cold fronts are frequently observed (Shafer and Steenburgh 2008). Further investigation revealed that smaller, weaker fronts in this area are rarely assigned a probability $\geq 65\%$ by the CNN, causing them to be thrown out. We found that including orographic height as a predictor causes more fronts to be detected in the Intermountain West but has a negative impact on the climatology in other areas.

The CNN used herein has notable advantages over other frontal-detection methods, which are mostly NFA or manual. First, CNNs outperform NFA at pattern recognition, because they incorporate machine learning with computer-vision tools, such as convolutional and pooling layers, allowing them to leverage learned features at different scales. Second, manual analysis is subjective (the same rules are not applied consistently over time) and too labor-intensive for a 40-yr climatology. Third, the CNN used here outperforms that developed by L19. Deterministic labels cannot be directly compared, because this study produces 2D frontal zones while L19 produces 1D lines. However, raw CNN probabilities can be directly compared. On the validation period (2015–16), the Gerrity score increases from 0.614 to 0.683, the Peirce score from 0.620 to 0.682, the Heidke score from

0.108 to 0.136, and the accuracy from 0.753 to 0.784. This performance gain is caused primarily by switching from the North American Regional Reanalysis to ERA5 (even the worst Gerrity scores in Fig. 3 are generally ≥ 0.614) and secondarily by combining two vertical levels in the CNN (the best Gerrity scores in Fig. 3 generally occur for models trained with surface data plus 950, 900, or 850 hPa).

The main limitation of our method is dependence on human labels for the initial CNN training. Human labels are unavailable for many feature-detection problems, and they are expensive to create from scratch. If a long archive of human labels is not available, we recommend training with labels from several diverse NFA methods, so that the CNN does not overfit the biases of one. Also, as mentioned above, human labels are subjective, even when based on a consensus of multiple experts. One tendency we have noticed is that humans label fronts more liberally in the summer, which appears to have been replicated by the CNN (Figs. 7 and 8). We considered removing CNN fronts with weak properties (e.g., thermal gradient or advection strength below a threshold), equivalent to using a masking variable in NFA. However, we decided against this, since it would reintroduce the vicissitudes of noisy gradients that often plague NFA. Another source of caution is that baroclinic troughs (Sanders 2005) are often identified by humans as cold fronts (Sanders and Hoffman 2002).

Our climatology has many potential applications beyond those explored in this study. For example, the association of fronts with thunderstorms (Fawbush and Miller 1954; Miller 1959) and extreme precipitation (Catto and Pfahl 2013; Dowdy and Catto 2017; Schemm et al. 2018b) suggests an opportunity to explore the long-term relationships of these phenomena with fronts. Previous attribution studies have typically used NFA or manual analysis to identify relevant fronts. We believe that the climatology developed herein could be used to learn more about these relationships. Also, we believe that combining machine learning with a 3D definition of fronts, such as that used in Spensberger and Sprenger (2018), would allow more to be learned about fronts and their climatology.

Acknowledgments. The Copernicus Climate Change Service was used to obtain ERA5 reanalysis data. The majority of computing for this project was performed at the OU Supercomputing Center for Education and Research (OSCER) at the University of Oklahoma (OU). All plots in this paper were generated with matplotlib, version 3.1.1 (Hunter 2007). We also thank the Weather Prediction Center (WPC) for creating the bulletins used to train the CNN and Alan Robson at

WPC for providing archived bulletins. Analyzed data for the study are available online (http://www.mcgovern-fagg.org/idea/theses/lagerquist_phd/index.html). The code to reproduce the CNN and analyses described can be found online (https://github.com/thunderhoser/GeneralExam/tree/era5_branch). We are grateful for the detailed feedback and suggestions of the three reviewers of the paper.

APPENDIX

Neighborhood Evaluation

CNN-produced frontal zones (section 2d) are scored via neighborhood evaluation. Traditional evaluation is spatially agnostic, comparing predicted and actual values grid cell by grid cell. This unduly punishes predictions that are correct except for a small displacement. This is especially problematic for comparing CNN fronts with WPC fronts, since 1) frontal placement is subjective and 2) we use 2D frontal zones while the WPC uses 1D lines. Neighborhood evaluation alleviates this problem by allowing a small displacement. We use neighborhood evaluation to compute four scores: probability of detection (POD), false-alarm ratio (FAR), frequency bias, and critical success index (CSI). We define these scores [Eq. (A1), below] differently than usual, because in our setting there are no nonfrontal objects, thus no correct nulls. Brooks (2004) use the same scores to evaluate NWS tornado warnings, where each case is a segment of a tornado track; since there is no such thing as a nontornado track, they also have no correct nulls.

$$\left\{ \begin{array}{l} \text{POD} = \frac{a_A}{a_A + c} \\ \text{FAR} = \frac{b}{a_P + b} \\ \text{bias} = \frac{\text{POD}}{1 - \text{FAR}} \\ \text{CSI}^{-1} = \text{POD}^{-1} + (1 - \text{FAR})^{-1} - 1 \end{array} \right. \quad (\text{A1})$$

The variables a_A , a_P , b , and c are defined as follows, where “actual” and “predicted” mean WPC and CNN, respectively: a_A = number of actual-oriented true positives, or actual frontal grid points matched with a predicted grid point of the same type (warm or cold), a_P = number of prediction-oriented true positives, or predicted frontal grid points matched with an actual grid point of the same type, b = number of false positives, or predicted frontal grid points *not* matched with an actual grid point of the same type, and c = number of false negatives, or actual frontal grid points *not* matched with a predicted grid point of the same type.

REFERENCES

- Allen, J., A. Pezza, and M. Black, 2010: Explosive cyclogenesis: A global climatology comparing multiple reanalyses. *J. Climate*, **23**, 6468–6484, <https://doi.org/10.1175/2010JCLI3437.1>.
- , M. Tippett, and A. Sobel, 2015: Influence of the El Niño/Southern Oscillation on tornado and hail frequency in the United States. *Nat. Geosci.*, **8**, 278–283, <https://doi.org/10.1038/ngeo2385>.
- American Meteorological Society, 2012: Wet-bulb potential temperature. *Glossary of Meteorology*, Accessed 13 March 2018, http://glossary.ametsoc.org/wiki/Pseudo_wet-bulb_potential_temperature.
- Berry, G., C. Jakob, and M. Reeder, 2011a: Recent global trends in atmospheric fronts. *Geophys. Res. Lett.*, **38**, L21812, <https://doi.org/10.1029/2011GL049481>.
- , M. Reeder, and C. Jakob, 2011b: A global climatology of atmospheric fronts. *Geophys. Res. Lett.*, **38**, L04809, <https://doi.org/10.1029/2010GL046451>.
- Beucler, T., T. Abbott, T. Cronin, and M. Pritchard, 2019: Comparing convective self-aggregation in idealized models to observed moist static energy variability near the equator. *Geophys. Res. Lett.*, **46**, 10 589–10 598, <https://doi.org/10.1029/2019GL084130>.
- Brenowitz, N., and C. Bretherton, 2019: Spatially extended tests of a neural network parametrization trained by coarse-graining. *J. Adv. Model. Earth Syst.*, **11**, 2728–2744, <https://doi.org/10.1029/2019MS001711>.
- Brooks, H., 2004: Tornado-warning performance in the past and future: A perspective from signal detection theory. *Bull. Amer. Meteor. Soc.*, **85**, 837–844, <https://doi.org/10.1175/BAMS-85-6-837>.
- Camargo, S., A. Robertson, A. Barnston, and M. Ghil, 2008: Clustering of eastern North Pacific tropical cyclone tracks: ENSO and MJO effects. *Geochem. Geophys. Geosyst.*, **9**, Q06V05, <https://doi.org/10.1029/2007GC001861>.
- Catto, J. L., and S. Pfahl, 2013: The importance of fronts for extreme precipitation. *J. Geophys. Res. Atmos.*, **118**, 10 791–10 801, <https://doi.org/10.1002/jgrd.50852>.
- , E. Madonna, H. Joos, I. Rudeva, and I. Simmonds, 2015: Global relationship between fronts and warm conveyor belts and the impact on extreme precipitation. *J. Climate*, **28**, 8411–8429, <https://doi.org/10.1175/JCLI-D-15-0171.1>.
- Chollet, F., 2018: *Deep Learning with Python*. Manning, 384 pp.
- Clarke, L., and R. Renard, 1966: The U.S. Navy numerical frontal analysis scheme: Further development and a limited evaluation. *J. Appl. Meteor.*, **5**, 764–777, [https://doi.org/10.1175/1520-0450\(1966\)005<0764:TUSNNF>2.0.CO;2](https://doi.org/10.1175/1520-0450(1966)005<0764:TUSNNF>2.0.CO;2).
- Climate Prediction Center, 2019: Monthly ERSSTv5 (1981–2010 base period) Niño-3.4 (5°N–5°S, 170°–120°W). Accessed 23 August 2019, <https://www.cpc.ncep.noaa.gov/data/indices/ersst5.nino.mth.81-10.ascii>.
- Cook, A., L. Leslie, D. Parsons, and J. Schaefer, 2017: The impact of El Niño–Southern Oscillation (ENSO) on winter and early spring U.S. tornado outbreaks. *J. Appl. Meteor. Climatol.*, **56**, 2455–2478, <https://doi.org/10.1175/JAMC-D-16-0249.1>.
- Crisp, C. A., and J. M. Lewis, 1992: Return flow in the Gulf of Mexico. Part I: A classificatory approach with a global historical perspective. *J. Appl. Meteor.*, **31**, 868–881, [https://doi.org/10.1175/1520-0450\(1992\)031<0868:RFITGO>2.0.CO;2](https://doi.org/10.1175/1520-0450(1992)031<0868:RFITGO>2.0.CO;2).
- Davis, S., and K. Rosenlof, 2012: A multidagnostic intercomparison of tropical-width time series using reanalyses and satellite observations. *J. Climate*, **25**, 1061–1078, <https://doi.org/10.1175/JCLI-D-11-00127.1>.

- Dieleman, S., K. Willett, and J. Dambre, 2015: Rotation-invariant convolutional neural networks for galaxy morphology prediction. *Nature*, **450**, 1441–1459, <https://doi.org/10.1093/MNRAS/STV632>.
- Dowdy, A., and J. Catto, 2017: Extreme weather caused by concurrent cyclone, front and thunderstorm occurrences. *Sci. Rep.*, **7**, 40359, <https://doi.org/10.1038/SREP40359>.
- Efron, B., 1979: Bootstrap methods: Another look at the jackknife. *Ann. Stat.*, **7** (1), 1–26, <https://doi.org/10.1214/aos/1176344552>.
- Eichler, T., and W. Higgins, 2006: Climatology and ENSO-related variability of North American extratropical cyclone activity. *J. Climate*, **19**, 2076–2093, <https://doi.org/10.1175/JCLI3725.1>.
- Fawbush, E., and R. Miller, 1954: The types of airmasses in which North American tornadoes form. *Bull. Amer. Meteor. Soc.*, **35**, 154–165, <https://doi.org/10.1175/1520-0477-35.4.154>.
- Fukushima, K., 1980: Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biol. Cybern.*, **36**, 193–202, <https://doi.org/10.1007/BF00344251>.
- , and S. Miyake, 1982: Neocognitron: A new algorithm for pattern recognition tolerant of deformations and shifts in position. *Pattern Recognit.*, **15**, 455–469, [https://doi.org/10.1016/0031-3203\(82\)90024-3](https://doi.org/10.1016/0031-3203(82)90024-3).
- Gagne, D., S. Haupt, and D. Nychka, 2019: Interpretable deep learning for spatial analysis of severe hailstorms. *Mon. Wea. Rev.*, **147**, 2827–2845, <https://doi.org/10.1175/MWR-D-18-0316.1>.
- Garner, J., 2013: A study of synoptic-scale tornado regimes. *Electron. J. Severe Storms Meteor.*, **8** (3), 1–25, <https://www.spc.noaa.gov/publications/garner/synoptic.pdf>.
- Goodfellow, I., Y. Bengio, and A. Courville, 2016: *Deep Learning*. MIT Press, 775 pp.
- Gossart, A., S. Helsen, J. Lenaerts, S. Vanden Broucke, N. P. M. van Lipzig, and N. Souverijns, 2019: An evaluation of surface climatology in state-of-the-art reanalyses over the Antarctic Ice Sheet. *J. Climate*, **32**, 6899–6915, <https://doi.org/10.1175/JCLI-D-19-0030.1>.
- Hardy, J. W., and K. G. Henderson, 2003: Cold front variability in the southern United States and the influence of atmospheric teleconnection patterns. *Phys. Geogr.*, **24**, 120–137, <https://doi.org/10.2747/0272-3646.24.2.120>.
- Henry, W., 1979: Some aspects of the fate of cold fronts in the Gulf of Mexico. *Mon. Wea. Rev.*, **107**, 1078–1082, [https://doi.org/10.1175/1520-0493\(1979\)107<1078:SAOTFO>2.0.CO;2](https://doi.org/10.1175/1520-0493(1979)107<1078:SAOTFO>2.0.CO;2).
- Hersbach, H., and D. Dee, 2016: ERA5 reanalysis is in production. *ECMWF Newsletter*, No. 147, ECMWF, Reading, United Kingdom, 7, <http://www.ecmwf.int/sites/default/files/elibrary/2016/16299-newsletter-no147-spring-2016.pdf>.
- Hewson, T., 1998: Objective fronts. *Meteor. Appl.*, **5**, 37–65, <https://doi.org/10.1017/S1350482798000553>.
- Hinton, G., N. Srivastava, A. Krizhevsky I. Sutskever, and R. Salakhutdinov, 2012: Improving neural networks by preventing co-adaptation of feature detectors. 18 pp., <https://arxiv.org/pdf/1207.0580.pdf>.
- Hodges, K. I., R. W. Lee, and L. Bengtsson, 2011: A comparison of extratropical cyclones in recent reanalyses ERA-Interim, NASA MERRA, NCEP CFSR, and JRA-25. *J. Climate*, **24**, 4888–4906, <https://doi.org/10.1175/2011JCLI4097.1>.
- Holzman, B., 1937: Detailed characteristics of surface fronts. *Bull. Amer. Meteor. Soc.*, **18**, 155–159, <https://doi.org/10.1175/1520-0477-18.4.5.155>.
- Hope, P., and Coauthors, 2014: A comparison of automated methods of front recognition for climate studies: A case study in southwest western Australia. *Mon. Wea. Rev.*, **142**, 343–363, <https://doi.org/10.1175/MWR-D-12-00252.1>.
- Hunter, J., 2007: Matplotlib: A 2D graphics environment. *Comput. Sci. Eng.*, **9**, 90–95, <https://doi.org/10.1109/MCSE.2007.55>.
- Hussain, M., and I. Mamhud, 2019: pyMannKendall: A Python package for non parametric Mann Kendall family of trend tests. *J. Open Source Software*, **4**, 1556, <https://doi.org/10.21105/joss.01556>.
- Ioffe, S., and C. Szegedy, 2015: Batch normalization: Accelerating deep network training by reducing internal covariate shift. *Int. Conf. on Machine Learning*, Lille, France, International Machine Learning Society, <http://proceedings.mlr.press/v37/loff15.pdf>.
- Kang, S., and L. Polvani, 2011: The interannual relationship between the latitude of the eddy-driven jet and the edge of the Hadley cell. *J. Climate*, **24**, 563–568, <https://doi.org/10.1175/2010JCLI4077.1>.
- Karras, T., T. Aila, S. Laine, and J. Lehtinen, 2018: Progressive growing of GANs for improved quality, stability, and variation. <https://arxiv.org/abs/1710.10196>.
- Kendall, M., 1955: *Rank Correlation Methods*. 2nd ed. Charles Griffin, 196 pp.
- Kurth, T., and Coauthors, 2018: Exascale deep learning for climate analytics. *Int. Conf. for High Performance Computing, Networking, Storage, and Analysis*, Dallas, TX, IEEE, <https://dl.acm.org/doi/pdf/10.5555/3291656.3291724>.
- Lagerquist, R., A. McGovern, and D. Gagne II, 2019: Deep learning for spatially explicit prediction of synoptic-scale fronts. *Wea. Forecasting*, **34**, 1137–1160, <https://doi.org/10.1175/WAF-D-18-0183.1>.
- Lee, S.-K., P. N. DiNezio, E.-S. Chung, S.-W. Yeh, A. T. Wittenberg, and C. Wang, 2014: Spring persistence, transition, and resurgence of El Niño. *Geophys. Res. Lett.*, **41**, 8578–8585, <https://doi.org/10.1002/2014GL062484>.
- Lucas, C., B. Timbal, and H. Nguyen, 2014: The expanding tropics: A critical assessment of the observational and modeling studies. *Wiley Interdiscip. Rev.: Climate Change*, **5**, 89–112, <https://doi.org/10.1002/wcc.251>.
- Maas, A., A. Hannun, and A. Ng, 2013: Rectifier nonlinearities improve neural network acoustic models. *Int. Conf. on Machine Learning*, Atlanta, GA, International Machine Learning Society, https://ai.stanford.edu/~amaas/papers/relu_hybrid_icml2013_final.pdf.
- Mann, H., 1945: Nonparametric tests against trend. *Econometrica*, **13**, 245–259, <https://doi.org/10.2307/1907187>.
- McGovern, A., R. Lagerquist, D. Gagne II, G. Jergensen, K. Elmore, C. Homeyer, and T. Smith, 2019: Making the black box more transparent: Understanding the physical implications of machine learning. *Bull. Amer. Meteor. Soc.*, **100**, 2175–2199, <https://doi.org/10.1175/BAMS-D-18-0195.1>.
- Metz, C., 1978: Basic principles of ROC analysis. *Semin. Nucl. Med.*, **8**, 283–298, [https://doi.org/10.1016/S0001-2998\(78\)80014-2](https://doi.org/10.1016/S0001-2998(78)80014-2).
- Miller, R., 1959: Tornado-producing synoptic patterns. *Bull. Amer. Meteor. Soc.*, **40**, 465–472, <https://doi.org/10.1175/1520-0477-40.9.465>.
- Morgan, G., D. Brunkow, and R. Beebe, 1975: Climatology of surface fronts. Illinois State Water Survey Circular 122 (ISWS-75-CIR122), 46 pp.
- Neu, U., and Coauthors, 2013: IMILAST: A community effort to intercompare extratropical cyclone detection and tracking algorithms. *Bull. Amer. Meteor. Soc.*, **94**, 529–547, <https://doi.org/10.1175/BAMS-D-11-00154.1>.
- Orlanski, I., 1975: A rational subdivision of scales for atmospheric processes. *Bull. Amer. Meteor. Soc.*, **56**, 527–530.
- Parfitt, R., A. Czaja, and H. Seo, 2017: A simple diagnostic for the detection of atmospheric fronts. *Geophys. Res. Lett.*, **44**, 4351–4358, <https://doi.org/10.1002/2017GL073662>.
- Payer, M., N. Laird, R. Maliawco, and E. Hoffman, 2011: Surface fronts, troughs, and baroclinic zones in the Great Lakes region. *Wea. Forecasting*, **26**, 555–563, <https://doi.org/10.1175/WAF-D-10-05018.1>.

- Racah, E., C. Beckham, T. Maharaj, and S. Kahou, Prabhat, and C. Pal, 2017: ExtremeWeather: A large-scale climate dataset for semi-supervised detection, localization, and understanding of extreme weather events. *Advances in Neural Information Processing Systems*, Long Beach, CA, Neural Information Processing Systems, <https://papers.nips.cc/paper/6932-extremeweather-a-large-scale-climate-dataset-for-semi-supervised-detection-localization-and-understanding-of-extreme-weather-events>.
- Rackley, J., and J. Knox, 2016: A climatology of southern Appalachian cold-air damming. *Wea. Forecasting*, **31**, 419–432, <https://doi.org/10.1175/WAF-D-15-0049.1>.
- Rakhlin, A., A. Shvets, V. Iglovikov, and A. Kalinin, 2018: Deep convolutional neural networks for breast cancer histology image analysis. <https://arxiv.org/abs/1802.00752>.
- Rasp, S., M. Pritchard, and P. Gentile, 2018: Deep learning to represent subgrid processes in climate models. *Proc. Natl. Acad. Sci. USA*, **115**, 9684–9689, <https://doi.org/10.1073/pnas.1810286115>.
- Renard, R., and L. Clarke, 1965: Experiments in numerical objective frontal analysis. *Mon. Wea. Rev.*, **93**, 547–556, [https://doi.org/10.1175/1520-0493\(1965\)093<0547:EINOFA>2.3.CO;2](https://doi.org/10.1175/1520-0493(1965)093<0547:EINOFA>2.3.CO;2).
- Roebber, P., 2009: Visualizing multiple measures of forecast quality. *Wea. Forecasting*, **24**, 601–608, <https://doi.org/10.1175/2008WAF2222159.1>.
- Rudeva, I., and I. Simmonds, 2015: Variability and trends of global atmospheric frontal activity and links with large-scale modes of variability. *J. Climate*, **28**, 3311–3330, <https://doi.org/10.1175/JCLI-D-14-00458.1>.
- , —, D. Crock, and G. Boschat, 2019: Midlatitude fronts and variability in the Southern Hemisphere tropical width. *J. Climate*, **32**, 8243–8260, <https://doi.org/10.1175/JCLI-D-18-0782.1>.
- Sanders, F., 1999: A proposed method of surface map analysis. *Mon. Wea. Rev.*, **127**, 945–955, [https://doi.org/10.1175/1520-0493\(1999\)127<0945:APMOSM>2.0.CO;2](https://doi.org/10.1175/1520-0493(1999)127<0945:APMOSM>2.0.CO;2).
- , 2005: Real front or baroclinic trough? *Wea. Forecasting*, **20**, 647–651, <https://doi.org/10.1175/WAF846.1>.
- , and C. A. Doswell, 1995: A case for detailed surface analysis. *Bull. Amer. Meteor. Soc.*, **76**, 505–522, [https://doi.org/10.1175/1520-0477\(1995\)076<0505:ACFDSA>2.0.CO;2](https://doi.org/10.1175/1520-0477(1995)076<0505:ACFDSA>2.0.CO;2).
- , and E. Hoffman, 2002: A climatology of surface baroclinic zones. *Wea. Forecasting*, **17**, 774–782, [https://doi.org/10.1175/1520-0434\(2002\)017<0774:ACOSBZ>2.0.CO;2](https://doi.org/10.1175/1520-0434(2002)017<0774:ACOSBZ>2.0.CO;2).
- Schemm, S., and M. Sprenger, 2015: Frontal-wave cyclogenesis in the North Atlantic—A climatological characterisation. *Quart. J. Roy. Meteor. Soc.*, **141**, 2989–3005, <https://doi.org/10.1002/qj.2584>.
- , I. Rudeva, and I. Simmonds, 2015: Extratropical fronts in the lower troposphere—global perspectives obtained from two automated methods. *Quart. J. Roy. Meteor. Soc.*, **141**, 1686–1698, <https://doi.org/10.1002/qj.2471>.
- , G. Rivière, L. M. Ciasto, and C. Li, 2018a: Extratropical cyclogenesis changes in connection with tropospheric ENSO teleconnections to the North Atlantic: Role of stationary and transient waves. *J. Atmos. Sci.*, **75**, 3943–3964, <https://doi.org/10.1175/JAS-D-17-0340.1>.
- , M. Sprenger, and H. Wernli, 2018b: When during their life cycle are extratropical cyclones attended by fronts? *Bull. Amer. Meteor. Soc.*, **99**, 149–165, <https://doi.org/10.1175/BAMS-D-16-0261.1>.
- Schmidt, D., and K. Grise, 2017: The response of local precipitation and sea level pressure to Hadley cell expansion. *Geophys. Res. Lett.*, **44**, 10 573–10 582, <https://doi.org/10.1002/2017GL075380>.
- Seager, R., N. Narnik, Y. Kushiner, W. Robinson, and J. Miller, 2003: Mechanisms of hemispherically symmetric climate variability. *J. Climate*, **16**, 2960–2978, [https://doi.org/10.1175/1520-0442\(2003\)016<2960:MOHSCV>2.0.CO;2](https://doi.org/10.1175/1520-0442(2003)016<2960:MOHSCV>2.0.CO;2).
- Sen, P., 1968: Estimates of the regression coefficient based on Kendall's tau. *J. Amer. Stat. Assoc.*, **63**, 1379–1389, <https://doi.org/10.1080/01621459.1968.10480934>.
- Serreze, M., and R. Barry, 2011: Processes and impacts of Arctic amplification: A research synthesis. *Global Planet. Change*, **77**, 85–96, <https://doi.org/10.1016/j.gloplacha.2011.03.004>.
- , A. Lynch, and M. Clark, 2001: The Arctic frontal zone as seen in the NCEP–NCAR reanalysis. *J. Climate*, **14**, 1550–1567, [https://doi.org/10.1175/1520-0442\(2001\)014<1550:TAFZAS>2.0.CO;2](https://doi.org/10.1175/1520-0442(2001)014<1550:TAFZAS>2.0.CO;2).
- Shafer, J., and W. Steenburgh, 2008: Climatology of strong intermountain cold fronts. *Mon. Wea. Rev.*, **136**, 784–807, <https://doi.org/10.1175/2007MWR2136.1>.
- Shapiro, L., and G. Stockman, 2000: *Computer Vision*. Pearson, 609 pp.
- Silver, D., and Coauthors, 2016: Mastering the game of Go with deep neural networks and tree search. *Nature*, **529**, 484–489, <https://doi.org/10.1038/nature16961>.
- , and Coauthors, 2017: Mastering the game of Go without human knowledge. *Nature*, **550**, 354–359, <https://doi.org/10.1038/nature24270>.
- Simmonds, I., K. Keay, and J. Bye, 2012: Identification and climatology of Southern Hemisphere mobile fronts in a modern reanalysis. *J. Climate*, **25**, 1945–1962, <https://doi.org/10.1175/JCLI-D-11-00100.1>.
- Spensberger, C., and M. Sprenger, 2018: Beyond cold and warm: An objective classification for maritime midlatitude fronts. *Quart. J. Roy. Meteor. Soc.*, **144**, 261–277, <https://doi.org/10.1002/qj.3199>.
- Suwajanakorn, S., S. M. Seitz, and I. Kemelmacher-Shlizerman, 2017: Synthesizing Obama: Learning lip sync from audio. *ACM Trans. Graph.*, **36**, 95, <https://doi.org/10.1145/3072959.3073640>.
- Tandon, N., E. Gerber, A. Sobel, and L. Polvani, 2013: Understanding Hadley cell expansion versus contraction: Insights from simplified models and implications for recent observations. *J. Climate*, **26**, 4304–4321, <https://doi.org/10.1175/JCLI-D-12-00598.1>.
- Tarek, M., F. P. Brissette, and R. Arseneault, 2020: Evaluation of the ERA5 reanalysis as a potential reference dataset for hydrological modelling over North America. *Hydrol. Earth Syst. Sci.*, **24**, 2527–2544, <https://doi.org/10.5194/hess-24-2527-2020>.
- Theil, H., 1992: A rank-invariant method of linear and polynomial regression analysis. *Henri Theil's Contributions to Economics and Econometrics*, B. Raj and J. Koerts, Eds., Vol. 23, Springer, 345–381.
- Thomas, B., and J. Martin, 2007: A synoptic climatology and composite analysis of the Alberta Clipper. *Wea. Forecasting*, **22**, 315–333, <https://doi.org/10.1175/WAF982.1>.
- Thomas, C., and D. Schultz, 2019a: Global climatologies of fronts, airmass boundaries, and airstream boundaries: Why the definition of “front” matters. *Mon. Wea. Rev.*, **147**, 691–717, <https://doi.org/10.1175/MWR-D-18-0289.1>.
- , and —, 2019b: What are the best thermodynamic quantity and function to define a front in gridded model output? *Bull. Amer. Meteor. Soc.*, **100**, 873–895, <https://doi.org/10.1175/BAMS-D-18-0137.1>.
- Wang, L., K. Scott, L. Xu, and D. Clausi, 2016: Sea ice concentration estimation during melt from dual-pol SAR scenes using deep convolutional neural networks: A case study. *IEEE Trans. Geosci. Remote Sens.*, **54**, 4524–4533, <https://doi.org/10.1109/TGRS.2016.2543660>.
- Wilks, D., 2016: “The stippling shows statistically significant grid points”: How research results are routinely overstated and overinterpreted, and what to do about it. *Bull. Amer. Meteor. Soc.*, **97**, 2263–2273, <https://doi.org/10.1175/BAMS-D-15-00267.1>.
- Wood, K., and E. Ritchie, 2014: A 40-year climatology of extratropical transition in the eastern North Pacific. *J. Climate*, **27**, 5999–6015, <https://doi.org/10.1175/JCLI-D-13-00645.1>.