

# FrontFinder AI: Efficient Identification of Frontal Boundaries over the Continental United States and NOAA's Unified Surface Analysis Domain Using the UNET3+ Model Architecture

ANDREW D. JUSTIN,<sup>a</sup> AMY MCGOVERN,<sup>a,b</sup> AND JOHN T. ALLEN<sup>c</sup>

<sup>a</sup> School of Meteorology, University of Oklahoma, Norman, Oklahoma

<sup>b</sup> School of Computer Science, University of Oklahoma, Norman, Oklahoma

<sup>c</sup> Department of Earth and Atmospheric Sciences, Central Michigan University, Mount Pleasant, Michigan

(Manuscript received 8 May 2024, in final form 29 October 2024, accepted 12 December 2024)

**ABSTRACT:** FrontFinder artificial intelligence (AI) is a novel machine learning algorithm trained to detect cold, warm, stationary, and occluded fronts and drylines. Fronts are associated with many high-impact weather events around the globe. Frontal analysis is still primarily done by human forecasters, often implementing their own rules and criteria for determining front positions. Such techniques result in multiple solutions by different forecasters when given identical sets of data. Numerous studies have attempted to automate frontal analysis through numerical frontal analysis. In recent years, machine learning algorithms have gained more popularity in meteorology due to their ability to learn complex relationships. Our algorithm was able to reproduce three-quarters of forecaster-drawn fronts over CONUS and NOAA's unified surface analysis domain on independent testing datasets. We applied permutation studies, an explainable artificial intelligence method, to identify the importance of each variable for each front type. The permutation studies showed that the most “important” variables for detecting fronts are consistent with observed processes in the evolution of frontal boundaries. We applied the model to an extratropical cyclone over the central United States to see how the model handles the occlusion process, with results showing that the model can resolve the early stages of occluded fronts wrapping around cyclone centers. While our algorithm is not intended to replace human forecasters, the model can streamline operational workflows by providing efficient frontal boundary identification guidance. FrontFinder has been deployed operationally at NOAA's Weather Prediction Center.

**SIGNIFICANCE STATEMENT:** Frontal boundaries drive many high-impact weather events worldwide. Identification and classification of frontal boundaries is necessary to anticipate changing weather conditions; however, frontal analysis is still mainly performed by human forecasters, leaving room for subjective interpretations during the frontal analysis process. We have introduced a novel machine learning method that identifies cold, warm, stationary, and occluded fronts and drylines without the need for high-end computational resources. This algorithm can be used as a tool to expedite the frontal analysis process by ingesting real-time data in operational environments.

**KEYWORDS:** Cold fronts; Fronts; Drylines; Warm fronts; Operational forecasting; Deep learning

## 1. Introduction

Fronts are boundaries and density contrasts that separate two air masses (Bjerknes 1919; Thomas and Schultz 2019) and influence the weather experienced around the globe daily. For example, fronts are the leading cause of extreme precipitation events in the United States (Kunkel et al. 2012) and can influence tornado development (Maddox et al. 1980; Childs and Schumacher 2019). These hazards motivate forecasters to identify these boundaries with a reasonable level of accuracy to better prepare end users for future weather events that pose threats to life and property. Forecasters at the National Weather Service (NWS) draw fronts every 3 h over North America and every 6 h over the unified surface analysis domain (USAD) which stretches from 130°E eastward to 10°E and from the equator northward to 80°N. Even among trained forecasters, the subjective nature of drawing fronts leads to a wide variety of solutions for the same data, as

illustrated by Uccellini et al. (1992). There is also a significant workload associated with drawing fronts over the entire USAD. In this paper, we present an improved machine learning (ML) algorithm for automated front detection that outperforms our previous models presented in Justin et al. (2023, hereafter J23).

Numerous studies have attempted to automate the frontal analysis process through gradients in selecting thermodynamic variables or wind shifts (e.g., Renard and Clarke 1965; Clarke and Renard 1966; Hewson 1998; Berry et al. 2011; Simmonds et al. 2012; Schemm et al. 2015). While these methods decrease subjectivity by removing the need for human forecasters to manually locate the fronts, the selection of parameters and rules for the various frontal analysis methods is a subjective process and yields different results with each method. The results are often noisy and not entirely focused on thermodynamic hypergradients. This has motivated research into the use of ML algorithms to further improve frontal analysis.

In recent years, the prevalence of ML in meteorological research papers has exhibited exponential growth (Chase et al.

*Corresponding author:* Andrew D. Justin, andrewjustin@ou.edu

DOI: 10.1175/AIES-D-24-0043.1 e240043

© 2025 American Meteorological Society. This published article is licensed under the terms of the default AMS reuse license. For information regarding reuse of this content and general copyright information, consult the AMS Copyright Policy ([www.ametsoc.org/PUBSReuseLicenses](http://www.ametsoc.org/PUBSReuseLicenses)).

## UNET3+ Architecture

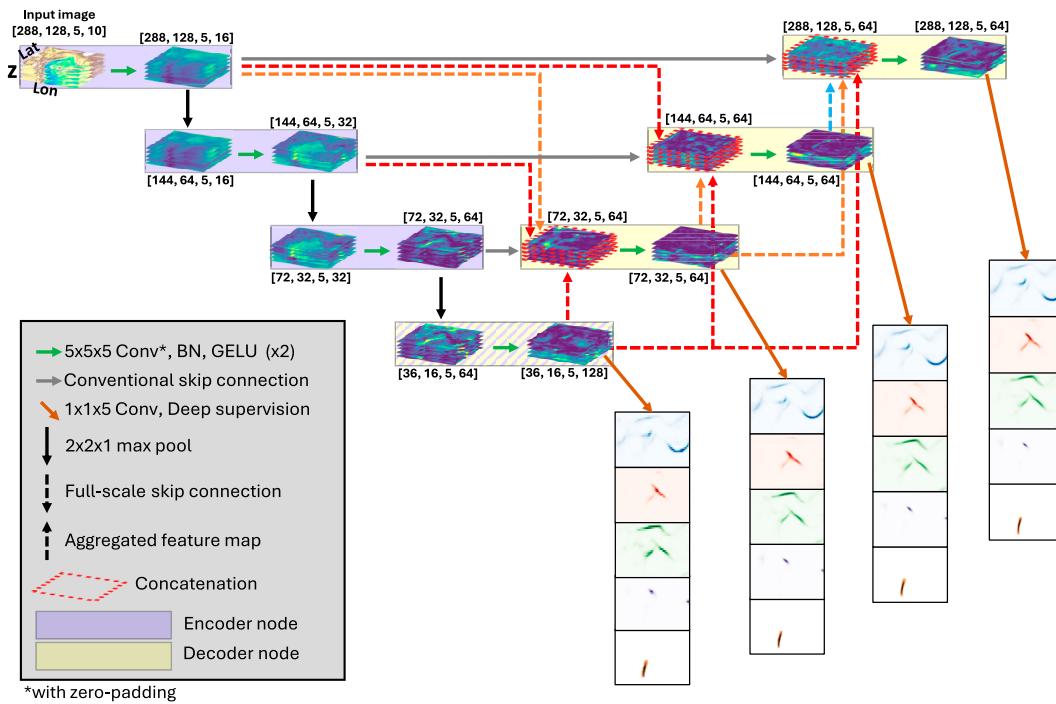


FIG. 1. Architecture of the UNET3+ model used to predict cold, warm, stationary, and occluded fronts and drylines. This example shows an input size of  $288 \times 128 \times 5 \times 10$ , where the third (vertical) dimension is unmodified until just prior to deep supervision. Note that the bottom node is both encoder and decoder nodes.

2022), and ML has shown great promise with automated front detection owing to its basis in the detection of sharp gradients. Lagerquist et al. (2019) used a convolutional neural network (CNN; LeCun et al. 1989) to locate cold and warm fronts over the domain covered by NWS' Weather Prediction Center (WPC), encompassing almost all of North America and surrounding areas. Their CNN drastically outperformed a baseline algorithm utilizing the thermal front parameter from Renard and Clarke (1965); however, the CNN was computationally expensive as it only made predictions for one pixel at a time. Biard and Kunkel (2019) used a 2D CNN trained on cold, warm, stationary, and occluded fronts from an archive of WPC's coded surface bulletin (CSB; National Weather Service 2019), achieving good categorical accuracy over WPC's domain. Lagerquist et al. (2020) trained multiple CNNs to identify and create climatologies of cold and warm fronts over North America, finding that the best-performing CNNs used a combination of input variables at the surface and 900 or 850 hPa. Bochenek et al. (2021) trained a random forest to predict fronts over Europe, using fronts drawn by analysts from the Deutscher Wetterdienst (DWD).

An alternate and more efficient approach is to use a specialized CNN architecture like the UNET (Ronneberger et al. 2015), designed for image segmentation and recognition (Denker and Burges 1995), which is able to make pixel-based predictions across an entire image. Matsuoka et al. (2019) used 2D UNET architectures to detect stationary fronts near

Japan, with some architectures only using one variable at a single pressure level as an input. Niebler et al. (2022) trained UNETs over North America and Europe and the North Atlantic, using fronts from the CSB and the DWD, respectively. They found that the models performed best over the domain on which they were trained (e.g., UNET trained on NWS data performed better over the NWS domain than the DWD), and models trained over both domains performed better overall. While applications of ML to frontal detection have been frequent, applications to the dryline have not (Clark et al. 2015).

TABLE 1. Predictor variables used as input to FrontFinder. Variables or levels in boldface were derived from variables directly retrieved from ERA5, which are not in boldface.

Variable	Heights/levels
Dewpoint temperature	2 m AGL, <b>1000, 950, 900, 850</b> hPa
Geopotential height	1000, 950, 900, 850 hPa
<b>Mixing ratio</b>	<b>2 m AGL, 1000, 950, 900, 850</b> hPa
<b>Relative humidity</b>	<b>10 m AGL, 1000, 950, 900, 850</b> hPa
Specific humidity	<b>2 m AGL</b> , 1000, 950, 900, 850 hPa
Surface pressure	Surface
Temperature	2 m AGL, 1000, 950, 900, 850 hPa
<b>Theta-E</b>	<b>2 m AGL, 1000, 950, 900, 850</b> hPa
U wind	10 m AGL, 1000, 950, 900, 850 hPa
V wind	10 m AGL, 1000, 950, 900, 850 hPa
<b>Virtual temperature</b>	<b>2 m AGL, 1000, 950, 900, 850</b> hPa

TABLE 2. Fraction of pixels containing each front type in the training, validation, and testing datasets. CF = cold front, WF = warm front, SF = stationary front, OF = occluded front, DL = dryline.

Dataset (years)	Image pairs	CF (%)	WF (%)	SF (%)	OF (%)	DL (%)
Training (2008–17, 2020)	34 766	1.751	0.523	1.491	0.222	0.0624
Validation (2018)	3232	1.954	0.608	1.81	0.193	0.075
Testing—CONUS (2019)	2920	2.027	0.65	1.67	0.238	0.0478
Testing—USAD (2019)	1460	1.009	0.377	0.772	0.291	0.0056

Clark et al. (2015) developed an automated dryline detection algorithm and showed that a random forest was able to detect more drylines than an algorithm not utilizing ML methods, while also lowering the false alarm ratio. To our knowledge, Clark et al. (2015) is the only study that has applied ML to dryline detection, leaving room for further exploration into developing ML-based algorithms that can efficiently detect drylines.

Our previous work (J23) used three sets of UNET3+ models (Huang et al. 2020) to predict cold, warm, stationary, and occluded fronts. We trained the models over CONUS and achieved good skill over both the CONUS and the USAD. However, there were three major downfalls with this old system of models, which will be referred to as the “three-model” system hereafter. First, the models were split into three sets: one for cold and warm fronts, one for stationary and occluded fronts, and another predicted the location of any of the four front types through binary classification (i.e., predicted the probability of any front type being present). Having a three-model system as opposed to a single model that predicts the front types means more time will be spent generating predictions, delaying the delivery of any real-time model outputs in operational settings. Second, the models in the three-model system were exceptionally large (roughly 233 million parameters each) and computationally inefficient. Finally, we were concerned with fronts being double counted and predicted given that the algorithms were trained independently, negatively affecting the overall skill or potentially leading to overpredictions. We also wanted to fill the gap with research into dryline detection, given that Clark et al. (2015) is the only study that has applied ML to dryline detection. To address these concerns, here we developed a new UNET3+ algorithm that includes the four classes of fronts along with drylines and illustrates how combining all front types into one algorithm

improves the algorithm’s ability to discern between the various front types while improving computational efficiency. This work is based on the primary author’s master’s thesis (Justin 2024).

## 2. Data and methods

### a. UNET3+ architecture

The UNET3+ is a CNN designed for image segmentation (Huang et al. 2020), improving upon the original UNET from Ronneberger et al. (2015) by adding “full-scale skip connections” and “aggregated feature maps” (more on this later in the section). In this case, the UNET3+ is assigning a class (front type) to every pixel of the model output. The architecture for our new UNET3+ model is shown in Fig. 1.

FrontFinder ingests 4D images as input, with the dimensions representing longitude, latitude, vertical level, and the predictor variables used in the model. In other words, multiple 3D variable maps are input to the model, with the variable maps stacked along the fourth dimension of the input images. Unlike our previous three-model system (J23), the longitude and latitude dimensions of FrontFinder are not fixed, and these dimensions of the input images can be changed under the condition that the lengths of the dimensions are evenly divisible by 16. In each encoder and decoder node of FrontFinder, the images/features pass through two convolution “modules,” with each module consisting of a  $5 \times 5 \times 5$  convolution operation, a batch normalization layer (Ioffe and Szegedy 2015), and a Gaussian error linear unit [GELU; Hendrycks and Gimpel (2023)] activation function layer. All  $5 \times 5 \times 5$  convolution operations throughout the model implement zero padding, a process in which layers of zeros are added around an image so that the output of the convolution operation has the same shape as the input (O’Shea and Nash 2015). Batch normalization layers normalize inputs such that the outputs have a mean of 0 and a variance of 1, which can provide numerical stability to the model (Ioffe and Szegedy

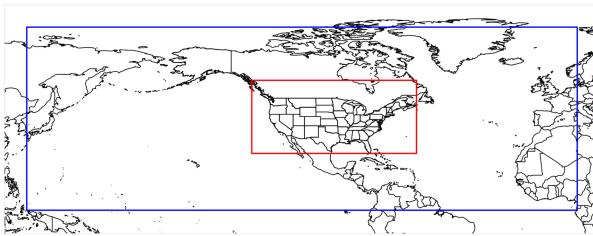


FIG. 2. CONUS domain (red) and NOAA’s USAD (blue). The bounds of the CONUS domain are  $25^{\circ}\text{N}$ ,  $56.75^{\circ}\text{N}$ ,  $132^{\circ}\text{W}$ ,  $60.25^{\circ}\text{W}$  ( $288 \times 128$  pixels on the  $0.25^{\circ}\text{ERA5}$  grid), while the USAD has bounds of  $0^{\circ}$ ,  $80^{\circ}\text{N}$ ,  $130^{\circ}\text{E}$ ,  $10^{\circ}\text{E}$  ( $960 \times 320$  pixels after 1-pixel truncation along each dimension).

TABLE 3. Hyperparameter choices for the FrontFinder algorithm.

Hyperparameter	Value
Batch size	64
Early stopping	55 epochs
Learning rate	$10^{-4}$
Loss function	Fraction skill score (Roberts 2008)
Optimizer	Adam (Kingma and Ba 2014)
Steps per epoch	Training: 10, validation: 51

TABLE 4. Performance of FrontFinder over (top) CONUS and (bottom) USAD. The three values in each cell represent scores using 50-, 100-, and 250-km neighborhoods. Drylines over USAD are omitted as they are not drawn outside WPC's domain.

Front type	CSI	POD (%)	FAR (%)	FB
CONUS				
Cold	0.507, 0.627, 0.682	70.9, 80.2, 86.3	36.1, 25.9, 23.6	1.110, 1.083, 1.130
Warm	0.311, 0.426, 0.505	54.8, 67.3, 78.1	58.2, 46.4, 41.2	1.310, 1.255, 1.329
Stationary	0.361, 0.458, 0.510	61.7, 74.0, 82.3	53.5, 45.4, 42.7	1.327, 1.354, 1.435
Occluded	0.386, 0.491, 0.541	61.2, 71.3, 76.2	48.9, 38.7, 34.9	1.199, 1.164, 1.170
Dryline	0.487, 0.624, 0.662	69.1, 78.6, 82.7	37.7, 24.8, 23.1	1.110, 1.046, 1.075
USAD				
Cold	0.466, 0.586, 0.645	66.8, 75.9, 80.9	39.3, 28.0, 23.9	1.102, 1.055, 1.062
Warm	0.313, 0.417, 0.480	54.0, 66.4, 75.1	57.3, 47.2, 42.9	1.264, 1.258, 1.315
Stationary	0.263, 0.339, 0.385	50.3, 61.8, 69.7	64.5, 57.1, 53.8	1.417, 1.440, 1.509
Occluded	0.353, 0.469, 0.540	56.0, 68.4, 74.6	51.2, 40.1, 33.8	1.148, 1.141, 1.128

2015). GELU was chosen as the activation function as it has been shown to provide superior performance over numerous other activation functions (Lee 2023). In our previous three-model system, we used the rectified linear unit (ReLU) activation function. ReLU and GELU are defined in Eqs. (1) and (2). ReLU outputs zero for all negative inputs; thus, the gradient or derivative for all negative inputs is also zero. If the inputs into a ReLU neuron are repeatedly negative, there is no gradient, rendering the neuron “dead” as its weights cannot be updated. This is known as the dying ReLU problem (Lu et al. 2020) and is one of the motivations of using GELU in place of ReLU. GELU is a smoother activation function that does not block out all negative inputs and allows gradients to flow smoothly through the model. ReLU and GELU are calculated as follows:

$$\text{ReLU}(x) = \begin{cases} x, & \text{if } x \geq 0 \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

$$\begin{aligned} \text{GELU}(x) &= x \times \frac{1}{2} [1 + \text{erf}(x/\sqrt{2})] \\ &\approx x \times \frac{1}{2} \{1 + \tanh[\sqrt{2/\pi}(x + 0.044715x^3)]\}. \end{aligned} \quad (2)$$

Maximum pooling operations connect two encoder nodes. Maximum pooling is a downsampling method that reduces the dimensions of an image using a predefined pool size, allowing UNET architectures to convert higher-resolution data into broader features that can be used by the model to further improve its performance. Since FrontFinder ingests 3D spatial images, the maximum pooling operations must be performed

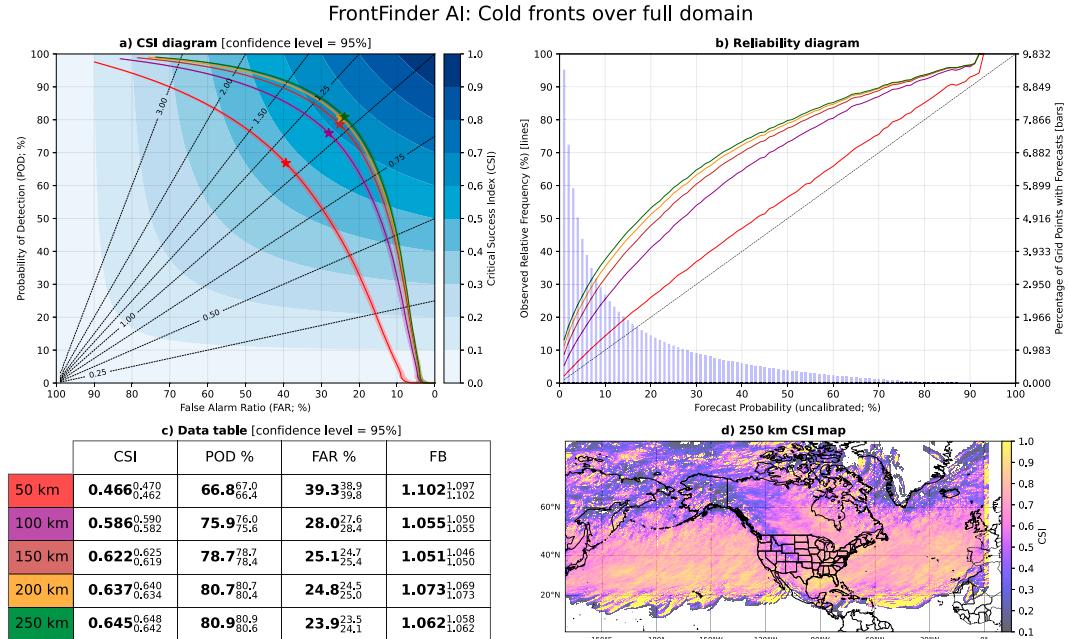


FIG. 3. Cold front results over USAD: (a) CSI diagram (dashed lines = FB), (b) reliability diagram (dashed line = perfect reliability), (c) table with upper and lower performance bounds indicated with superscripts and subscripts, respectively, and (d) spatial CSI diagram using a 250-km neighborhood. CSI scores lower than 0.1 are not shown on the spatial diagram.

## Cold front permutations: Unified Surface Analysis Domain

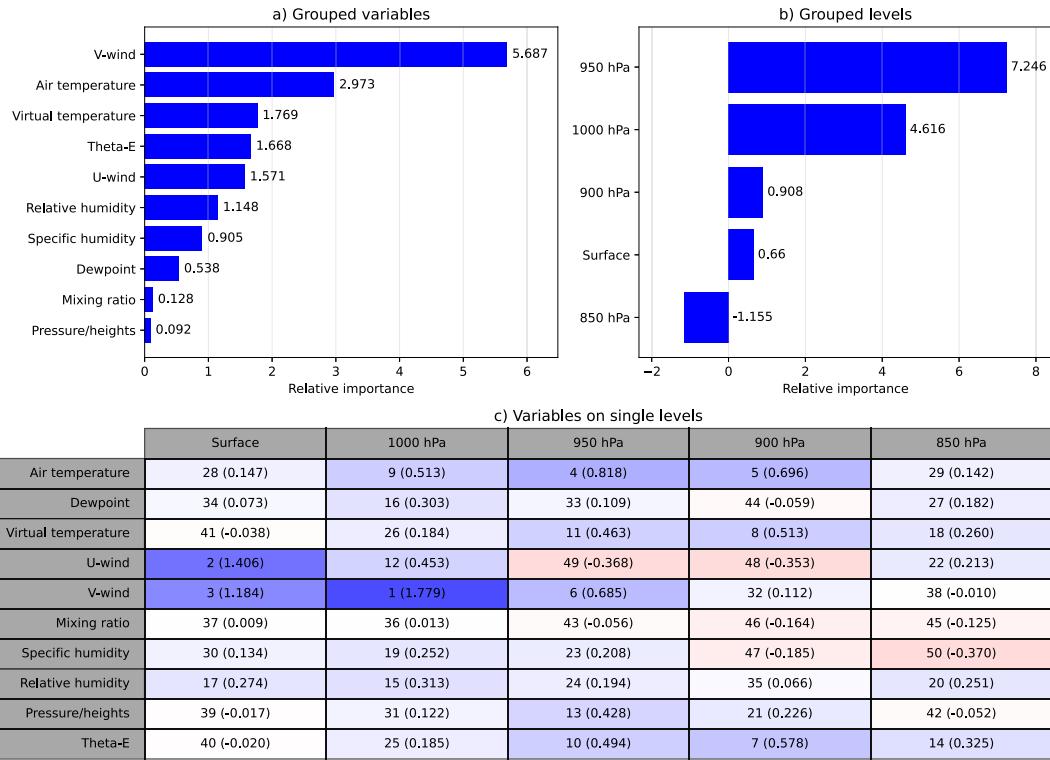


FIG. 4. Cold front permutation results over the USAD for (a) grouped variables, (b) grouped vertical levels, and (c) variables on single levels ranked from 1 to 60 with 1 (60) being the most (least) important variable and level combination.

with 3D pool sizes. However, the vertical dimension of our input images only has a size of 5, so shrinking the size of the vertical dimension with maximum pooling in our architecture is not feasible. We preserve the size of the vertical dimension while still reducing horizontal dimensions by using a pool size of  $2 \times 2 \times 1$  in our maximum pooling operations. We connect encoder and decoder nodes on the same level of FrontFinder with a conventional skip connection, which takes the output of an encoder node and passes it through one convolution module before connecting it to the decoder node. Full-scale skip connections transport data from a high-resolution encoder node to a lower-resolution decoder node; images undergo a 3D maximum pooling operation before passing through a convolution module and being transported to the decoder node. The pool size in each full-scale skip connection is determined by the resolutions of the images that are processed by each of the nodes attached to the connection. For example, if a full-scale skip connection sends images with dimensions  $288 \times 128 \times 5$  to a decoder node that processes images with dimensions  $72 \times 32 \times 5$ , then the pool size must be  $4 \times 4 \times 1$ . On the decoder/upsampling side of the UNET3+, nodes are connected by connections containing aggregated feature maps. Aggregated feature maps are low-resolution images that have undergone 3D upsampling operations and passed through one convolution module before reaching the target decoder node.

The decoder nodes create concatenated features from all incoming connections before passing the features through a convolution module and performing deep supervision. In FrontFinder, deep supervision in the decoder nodes (and the bottom encoder/decoder node) starts with features undergoing  $5 \times 5 \times 5$  convolutions with zero padding, followed by  $1 \times 1 \times 5$  convolutions without zero padding. The absence of zero padding means that the  $1 \times 1 \times 5$  convolutions shrink the vertical dimension from size 5 to size 1. We then “squeeze out” the vertical dimension since it has a size of 1, resulting in 2D feature maps with longitude and latitude dimensions. With the exception of the final decoder node, the 2D feature maps are upsampled with pool sizes that result in shapes matching the longitude and latitude dimensions of the original input images. After upsampling, the features go through a Softmax function (Bridle 1989) and turn into probabilities for all the front types in the model.

#### b. Datasets and preprocessing

Datasets were made with predictor variables sourced from ERA5 data (Hersbach et al. 2023) and front positions as analyzed by NOAA forecasters at 3-h intervals (NOAA 2023) for the period 2008–20. ERA5 data were chosen given its high spatial resolution ( $0.25^\circ$ , roughly 25 km) and global coverage with data at numerous pressure levels available every 3 h.

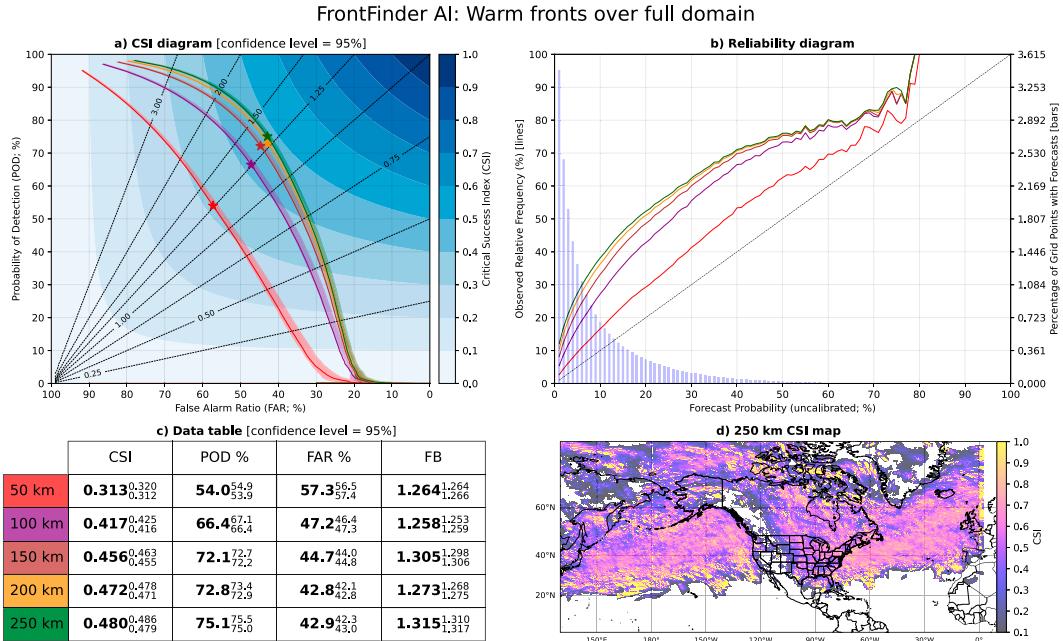


FIG. 5. Warm front results over USAD: (a) CSI diagram (dashed lines = FB), (b) reliability diagram (dashed line = perfect reliability), (c) table with upper and lower performance bounds indicated with superscripts and subscripts, respectively, and (d) spatial CSI diagram using a 250-km neighborhood. CSI scores lower than 0.1 are not shown on the spatial diagram.

ERA5 reanalysis has a fixed model state, so the data assimilation process is consistent with time. While the reanalysis data are quality controlled, they may have spatial biases in boundary positions and be unable to capture some extreme mesoscale thermodynamic gradients, like those seen with some drylines (Pietrycha and Rasmusson 2004). Using NOAA's fronts allows us to have a high temporal resolution with some consistency between subsequent analyses since the same offices are drawing the analyses. Ten predictor variables at five vertical levels were used as input (Table 1), while the labels were front positions interpolated every 1 km and transformed to a uniform  $0.25^\circ$  grid to match the ERA5 resolution. This differs from our previous three-model system as wet-bulb temperature and wet-bulb potential temperature  $\theta_W$  are not included in the list of predictors since our permutation studies in J23 suggested that these variables have little influence on the model predictions. The front labels were also expanded by one pixel ( $0.25^\circ$  or roughly 25 km) in all directions in order to account for positional biases that may exist between the reanalysis and observed boundaries. The data were split into three datasets—training, validation, and testing. The training dataset encompassed the years 2008–17 and 2020, while the validation and testing datasets contained data for 2018 and 2019, respectively. The discontinuity in the years used for the training set is intentional; 2020 had the highest sample of drylines throughout our entire dataset (Table 2), so it was included in the training data to maximize dryline sample sizes. The training and validation datasets only cover our CONUS domain. Note that the testing dataset for USAD only contains 6-hourly time steps (0z, 6z, 12z, 18z) as full analyses over

USAD are only performed every 6 h. The extents of the CONUS domain and USAD are shown in Fig. 2.

Time steps used in the training and validation datasets were selected with a two-step process. First, each 3-hourly time step in the years within each dataset is checked for the presence of at least one of each front type over the CONUS domain. If at least one of each type is present, the time step is retained in the dataset. Otherwise, the time step is retained with a 50% chance. Filtering is necessary to compensate for the large variation in sample sizes between the different front types. This process results in approximately 92% of the time steps in the datasets being randomly selected with no seasonality or time dependence, which subsequently limits any possible seasonal biases, especially with front types whose frequencies have a strong seasonal dependency (e.g., drylines in the spring). Nine pairs of images evenly spaced along the longitude are extracted from each time step; the shape of the inputs was  $128 \times 128 \times 5 \times 10$  (longitude  $\times$  latitude  $\times$  vertical level  $\times$  variable), while the labels containing fronts have shape  $128 \times 128 \times 6$  (longitude  $\times$  latitude  $\times$  front type; one of the front types is labeled “no front”). To prevent the model from learning spatial correlations that exist with some front types which may not be transferable outside of the training domain (e.g., the north–south orientation of drylines in the Great Plains with dry air to the west and moist air to the east of the boundary), each horizontal dimension in the pairs of images had a 25% chance of being flipped. In other words, 43.75% of all images had at least one horizontal dimension flipped. The process of increasing the size of a dataset with modified copies of images is known as data augmentation, which has been

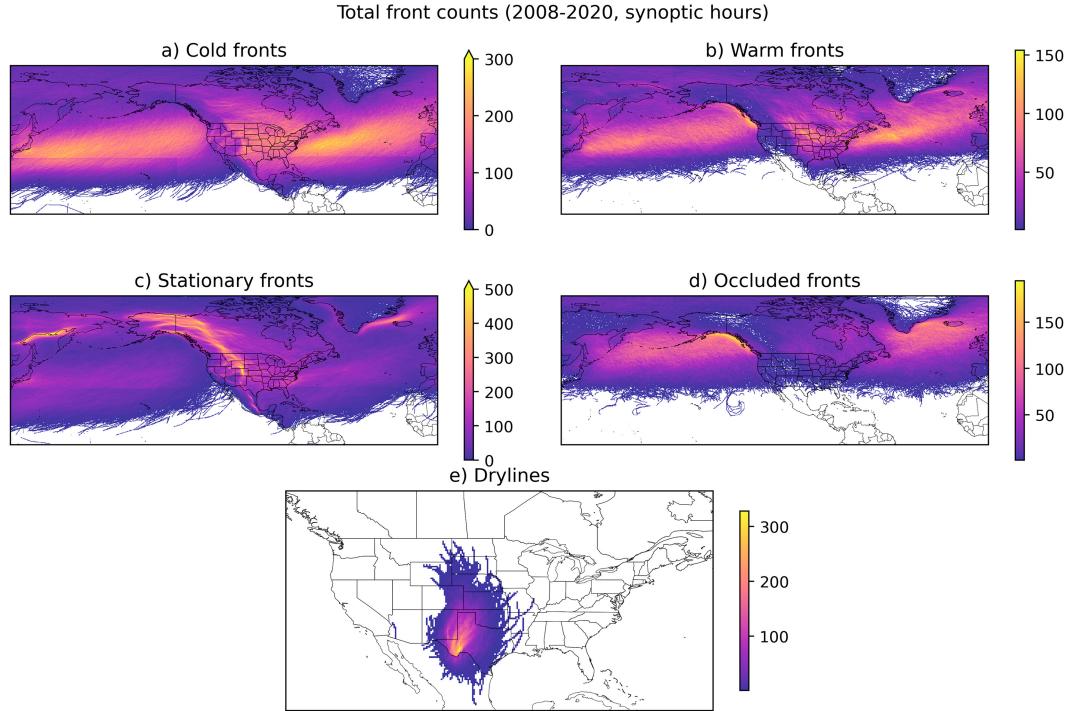


FIG. 6. Frequency of (a) cold fronts, (b) warm fronts, (c) stationary fronts, (d) occluded fronts, and (e) drylines drawn by NWS forecasters over USAAD for the period 2008–22 at synoptic hours. Nonsynoptic hours are not shown as WPC only draws over North America for nonsynoptic time steps and frequencies are significantly higher over the WPC domain (see Fig. 2 from J23).

shown to improve the performance of deep learning models applied to object recognition tasks (Yang et al. 2023). Note that the data were not augmented based on correlations between the input variables, which could have implications for results stemming from our permutation studies. For example, as moisture content increases, several variables will have their values increased (e.g., relative humidity, specific humidity, mixing ratio, and dewpoint temperature). Correlations between the variables may not assign importance based on their true predictive abilities (e.g., splitting importance up between two closely linked variables, like specific humidity and mixing ratio). Ablation studies were not performed to assess how correlations between the input variables impact model performance, so it is unclear if better performance can be achieved by removing these correlations from the inputs. Nonetheless, variables were normalized with minimum–maximum normalization at each vertical level to ensure predictor variables are in a common range of [0, 1] and provide numerical stability to the model.

#### c. Training and hyperparameters

FrontFinder was trained in parallel across four NVIDIA A100 graphics processing units (GPUs) (40-GB variant) over the course of 18.5 h. The hyperparameter choices are summarized in Table 3.

The loss function is a key part of the training process and serves as an error metric for the model and optimizes model

parameters (Terven et al. 2023). We used a simple manipulation of the fraction skill score [FSS; Roberts (2008)] as the loss function for the model. FSS is a spatial verification metric that does not penalize predictions that are displaced from the ground truth by a predefined number of pixels, where an FSS of 1 (0) is the best (worst) possible forecast. In the case of FrontFinder, predictions within one pixel ( $0.25^\circ$  or roughly 25 km) of a ground truth front are considered hits. The formula for our loss function  $FSS_{loss}$  is shown in Eq. (3) below, with  $FSS_{loss}$  equal to 0 (1) for the best (worst) possible forecast:

$$FSS_{loss} = 1 - FSS. \quad (3)$$

Adam (Kingma and Ba 2014) was used as the optimizer, with an initial learning rate of  $10^{-4}$  and using the default exponential decay rates from Kingma and Ba (2014). The batch size for both training and validation was 64, performing 10 steps for training and 51 steps for validation in each epoch. In our case, an epoch is defined by one pass over a small subset of the data (10 passes/steps over batches of 64 images). Validation is performed every epoch, and 51 steps ensure that the model sees all images in the validation dataset every time that validation is performed.

Early stopping was used to prevent the model from overfitting to the training dataset. This method of regularization involves stopping the training process early when model error

## Warm front permutations: Unified Surface Analysis Domain

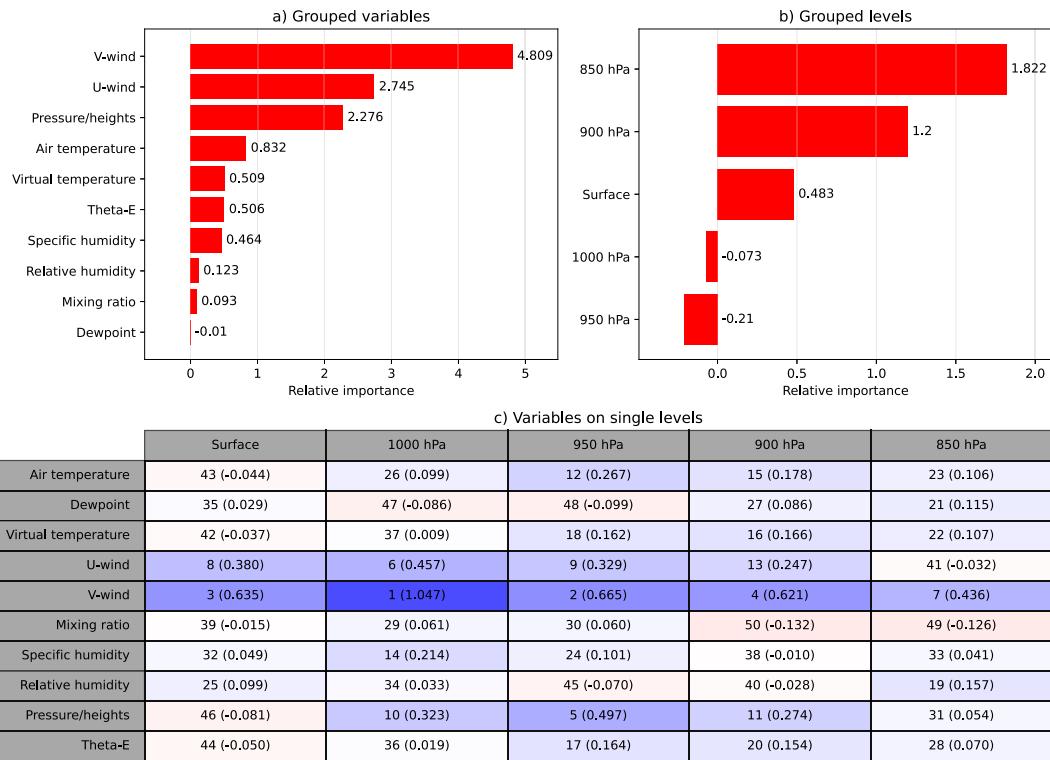


FIG. 7. Warm front permutation results over the USAD for (a) grouped variables, (b) grouped vertical levels, and (c) variables on single levels ranked from 1 to 60 with 1 (60) being the most (least) important variable and level combination.

fails to decrease, the exact conditions of which are often subjective and determined prior to model training [see [Shen et al. \(2022\)](#) for more details]. Training for the model was stopped after the validation loss failed to improve for 55 successive epochs. This rule was set to ensure that the model completed a full iteration over the training dataset without improvements before training was suspended. The model achieved the lowest validation loss at epoch 644, with training being terminated after epoch 699 following 55 consecutive epochs of no improvements in the validation loss.

#### d. Evaluation

To evaluate the performance of FrontFinder over the CONUS domain and USAD, we generated model predictions with images from the respective testing sets for each domain (see [Table 2](#)) and calculated the critical success index (CSI) using 50-, 100-, 150-, 200-, and 250-km neighborhoods (assuming 25 km between grid spaces). Similar to FSS, CSI evaluation using a neighborhood allows for predictions slightly displaced from the ground truth to be counted as hits (e.g., see [Clark et al. 2015; J23; Lagerquist et al. 2019, 2020; Niebler et al. 2022](#)). CSI is calculated at probability thresholds from 0.01 to 1 (0.01, 0.02, 0.03, ..., 1) and is defined in Eq. (4) as follows:

$$CSI = \frac{TP}{TP + FP + FN} = \frac{1}{\frac{1}{POD} + \frac{1}{1 - FAR} - 1}, \quad (4)$$

where TP, FP, and FN are the numbers of true positives, false positives, and false negatives, respectively; POD is the probability of detection [Eq. (5)]; and FAR is the false alarm ratio [Eq. (6)]:

$$POD = \frac{TP}{TP + FN}, \quad (5)$$

$$FAR = \frac{FP}{TP + FP}. \quad (6)$$

Frequency bias (FB) was also calculated in order to gauge the model's tendency to overpredict or underpredict certain types. FB is defined in Eq. (7) as follows:

$$FB = \frac{TP + FP}{TP + FN}. \quad (7)$$

When FB is greater (less) than 1, the model tends to overpredict (underpredict) events, so a value of 1 is preferred. Confidence intervals at the 95% level were calculated for bulk performance statistics using bootstrapping, iterating over the statistics 1000 times and retaining statistics for 2920 (1460)

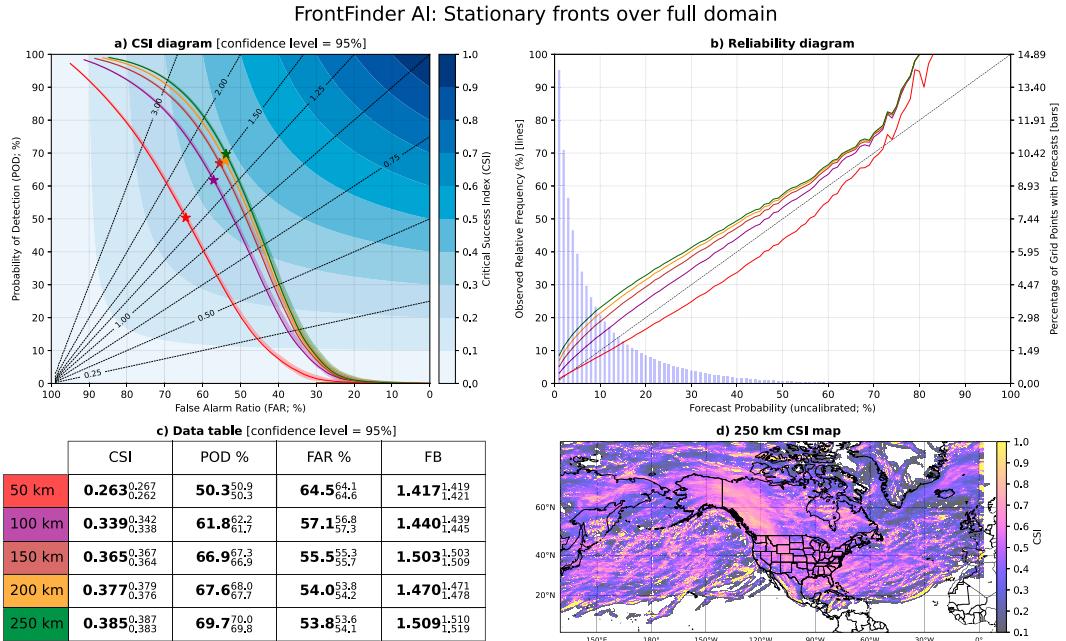


FIG. 8. Stationary front results over USAD: (a) CSI diagram (dashed lines = FB), (b) reliability diagram (dashed line = perfect reliability), (c) table with upper and lower performance bounds indicated with superscripts and subscripts, respectively, and (d) spatial CSI diagram using a 250-km neighborhood. CSI scores lower than 0.1 are not shown on the spatial diagram.

time steps each iteration when bootstrapping over the CONUS domain (USAD).

#### e. Permutation studies

Variable and pressure-level importance was determined through single-pass permutations (Breiman 2001; McGovern et al. 2019). Single-pass permutations involve shuffling values of a particular predictor to see how model performance changes. In our case, we shuffled each variable and/or pressure level in all images within the testing datasets and reevaluated the model across the new datasets containing the shuffled data. Sixty-five permutations were performed for each front type: 50 for single variables at individual pressure levels, 10 for single variables over all pressure levels, and five for all variables at individual pressure levels. The change in POD was used as a metric for importance. The decision to use POD over CSI for importance was made because fronts with exceptionally small sample sizes (e.g., drylines) can see the CSI actually increase during permutation studies due to model probabilities decreasing across the pixels without any targets (in the case of drylines, >99.95% of pixels in the CONUS testing set are empty). This can cause a misleading result whereby parameter importance is understated due to a large drop in the number of false positives. Using POD prevents this issue and only focuses on pixels containing the front types of interest. A decrease in POD indicates that a predictor is important and vice versa for an increase in POD. The change in POD was pixel based and used the dilated front labels as the targets without any neighborhood approaches like those used in our CSI calculations.

### 3. Results and discussion

We first evaluated FrontFinder's performance for each of the five front types with the testing dataset (2019) over CONUS and USAD (Table 4). Figures 3–11 cover USAD and are in the main body of this paper, while CONUS diagrams (except for drylines) are located in the appendixes A and B (Figs. A1–A4, B1–B4). Over CONUS, cold fronts achieved the best performance of all front types with a 250-km CSI score of 0.682, successfully hitting 86.3% of pixels containing cold fronts with a 23.6% FAR, outperforming our three-model system from J23 that achieved a CSI of 0.55 for cold fronts over CONUS. The high cold front performance from FrontFinder is consistent with findings from J23 and Niebler et al. (2022); both papers found that cold fronts were the best-performing boundary type over multiple domains, excluding the binary front type. Cold front performance is weaker over the Rocky Mountains and areas north of 60°N (Figs. 3 and A1d). The number of cold fronts predicted by the model was slightly higher than the number of ground truth fronts in the test dataset, leading to FB scores between 1.05 and 1.10 over USAD. By comparison, cold front FB scores in J23 over USAD were between 1.10 and 1.27, suggesting that the simplification of the model architecture reduced the tendency to overpredict cold fronts. Terrain and the lack of input data at higher altitudes likely complicate cold front detection over the Rocky Mountains. For example, mountains can complicate frontal structures by blocking and diverting flows, and cold air can sink into valleys and clash with existing air masses. The

## Stationary front permutations: Unified Surface Analysis Domain

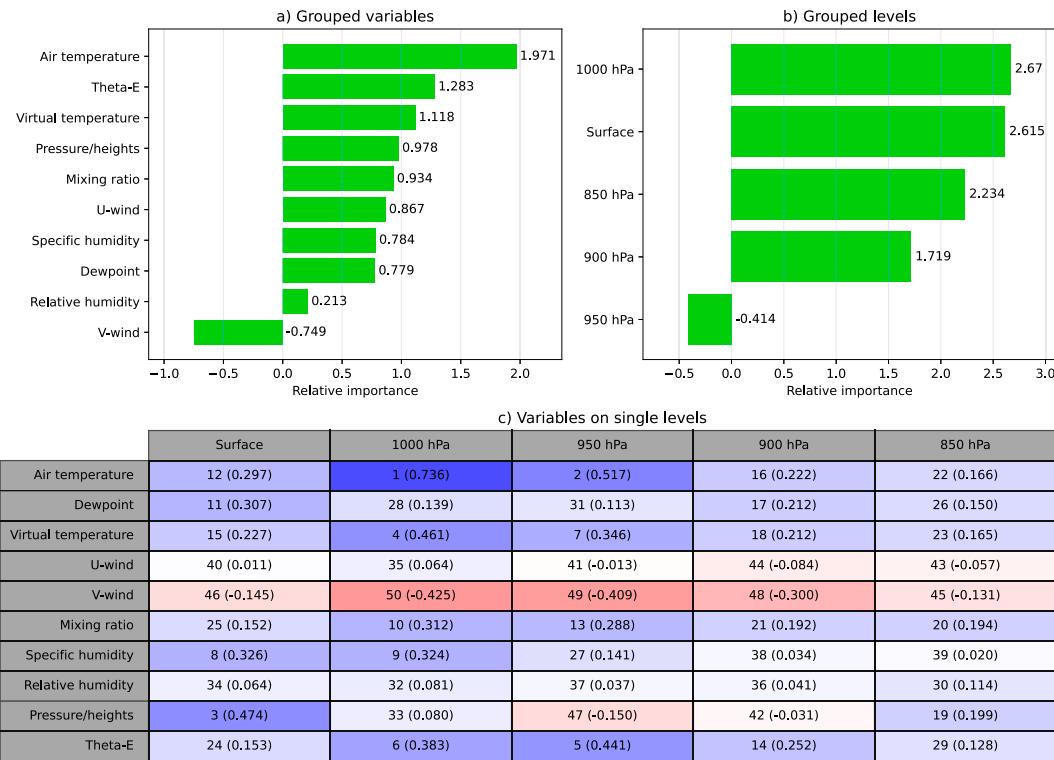


FIG. 9. Stationary front permutation results over the USAD for (a) grouped variables, (b) grouped vertical levels, and (c) variables on single levels ranked from 1 to 60 with 1 (60) being the most (least) important variable and level combination.

lower frequency of cold fronts at higher latitudes and the fact that the distance between longitudes decreases with northern extent could be an explanation for lower performance over higher latitudes (Lagerquist et al. 2019). In the midlatitudes, there is little difference in cold front performance between the CONUS and the Atlantic and Pacific Ocean basins, which indicates that FrontFinder is able to generalize oceanic cold fronts despite the training data only covering a very small portion of these basins.

The cold front permutation studies shown in Figs. 4 and B1 indicate that v wind, temperature, and virtual temperature are the most important variables for cold front detection. This is not a surprising result since cold fronts are characterized by temperature contrasts and often wind shifts and moisture gradients (Schultz 2005). One interesting result was that surface pressure and geopotential heights had much lower importance over USAD versus CONUS. We believe this is related to surface friction; overhanging noses of denser air ahead of surface boundaries have been documented in several studies (e.g., Simpson 1972; Young and Johnson 1984; Mitchell and Hovermale 1977), with Simpson (1972) attributing the overhanging nose to a no-slip lower boundary condition. Since USAD covers much of the Pacific and Atlantic Oceans, we think that the higher importance for surface pressure and

geopotential heights over CONUS is the result of overhanging noses caused by friction with land surfaces.

FrontFinder performance for warm fronts was similar across the CONUS domain and USAD with 250-km CSI scores of 0.505 and 0.48, respectively, a significant improvement over the three-model system in J23 that achieved CSI scores of 0.36 and 0.37. FrontFinder found 75.1% of the warm fronts within 250 km of the ground truth warm fronts as analyzed by forecasters with a FAR of 42.9% over USAD. FB scores for warm fronts over USAD ranged from 1.26 to 1.32 with FrontFinder and 1.26–1.49 with the three-model system in J23, suggesting that FrontFinder is more consistent with warm front predictions. CSI scores are slightly higher over water, in particular the Atlantic and Pacific Ocean basins, than over land (Figs. 5 and A2d). This could be attributed to differences in the intensity of warm fronts over water versus land; Hines and Mechoso (1993) analyzed the evolution of frontal structures during numerical simulations of cyclogenesis and found that simulations with low surface drag exhibited enhanced warm frontogenesis due to the higher wind speeds from low surface drag allowing for more robust warm air advection into the warm frontal zones. Warm fronts are rarely analyzed over the Rocky Mountains (Fig. 6b), which implies that sample size is the likely driver of lower CSI scores over the mountainous terrain. The simplified model

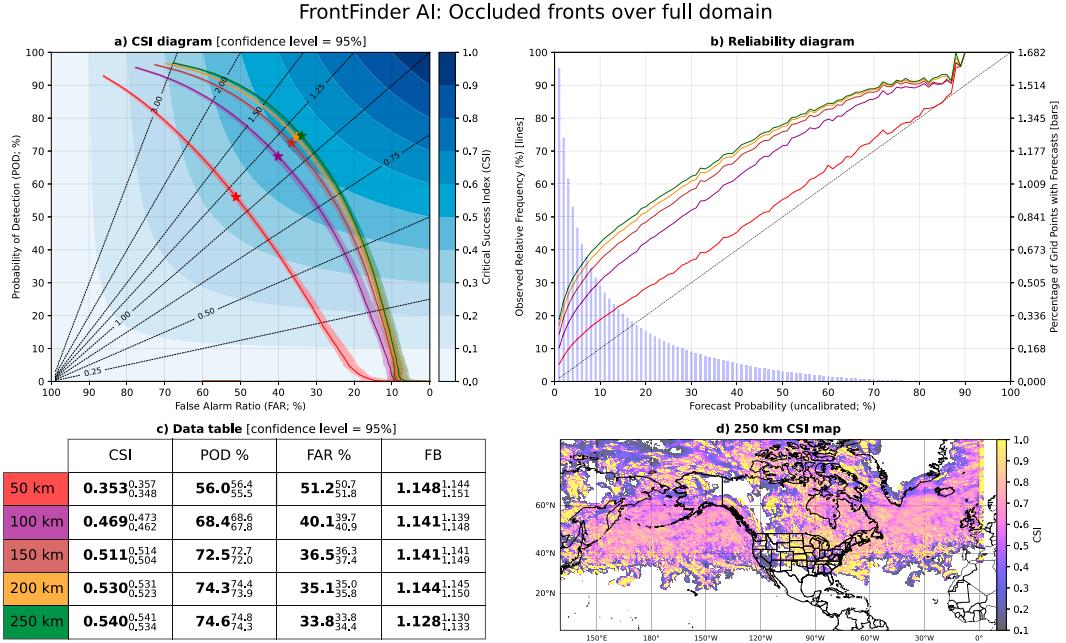


FIG. 10. Occluded front results over USAD: (a) CSI diagram (dashed lines = FB), (b) reliability diagram (dashed line = perfect reliability), (c) table with upper and lower performance bounds indicated with superscripts and subscripts, respectively, and (d) spatial CSI diagram using a 250-km neighborhood. CSI scores lower than 0.1 are not shown on the spatial diagram.

architecture used in the present study that includes all front types likely helped FrontFinder to better differentiate between warm and stationary fronts, both of which experienced high FAR in J23.

The warm front permutation studies shown in Figs. 7 and B2 echoed findings of an observation of a warm front over the northeast Atlantic by Wakimoto and Bosart (2001). They found that the observed warm front was more defined aloft with a stronger equivalent potential temperature  $\theta_E$  gradient and intense vertical wind shear. Wakimoto and Bosart (2001) also noted that the warm front at the surface had little  $\theta_E$  or mixing ratio gradient. Our permutation studies over USAD showed that the 850- and 900-hPa levels had greater relative importance than the surface, 1000, and 950 hPa. The results shown in Fig. 7c show  $\theta_E$  has the greatest importance at 950, 900, and 850 hPa over USAD, with the surface  $\theta_E$  having negative relative importance (i.e., hurting warm front predictability). Figure 7a shows that v wind and u wind were the most important variables in detecting warm fronts over USAD, which is supported by the observation of intense vertical wind shear along the oceanic warm front analyzed by Wakimoto and Bosart (2001). The permutation studies suggest that surface pressure is not important to detecting warm fronts; however, this may be due to emphasis on geopotential heights of pressure levels above the surface.

Stationary front performance is among FrontFinder's weaknesses, with the model only achieving moderate improvement over J23, where CSI scores were 0.44 and 0.27. A local maximum in CSI can be noted along the Rocky Mountains in Alberta and British Columbia (Figs. 8 and A3d), coincident with

a higher frequency of stationary fronts (Fig. 6c). While stationary front performance moderately improved over J23, stationary fronts still have the highest FAR of any front type in the model. Exploring over the unified domain, we found that FrontFinder tends to falsely identify parts of the intertropical convergence zone (ITCZ) as a stationary front. The ITCZ is not included in the training data, so we think that the model has not learned how to properly interpret semipermanent convergence zones. Stationary fronts seem to be one of the more subjective parts of surface analysis, and the transition of a cold or warm front to a stationary front is not always definitive. This implies a potential issue with other quasi-permanent boundary features and the approach to training on individual time steps, both of which could be sources of the high FAR for stationary fronts. The model only ingests data for one time step, so the model has no information on previous positions of fronts and is unable to base predictions on the movement of the fronts.

The stationary front permutation studies (Figs. 9 and B3) indicate that air temperature is the most important variable for stationary front detection. V wind has a relatively strong negative effect on FrontFinder's ability to detect stationary fronts; however, the reason for this is not clear. U wind has stronger importance over the CONUS domain compared to USAD, which we suspect to be linked to the predominantly zonal orientation of stationary fronts frequently observed along and over the Rocky Mountains (Fig. 6c). We also expected greater importance over CONUS for pressure levels at greater elevation due to the elevated and complex terrain over the Rocky Mountains; however, only the 950-hPa level had notably higher importance over CONUS.

## Occluded front permutations: Unified Surface Analysis Domain

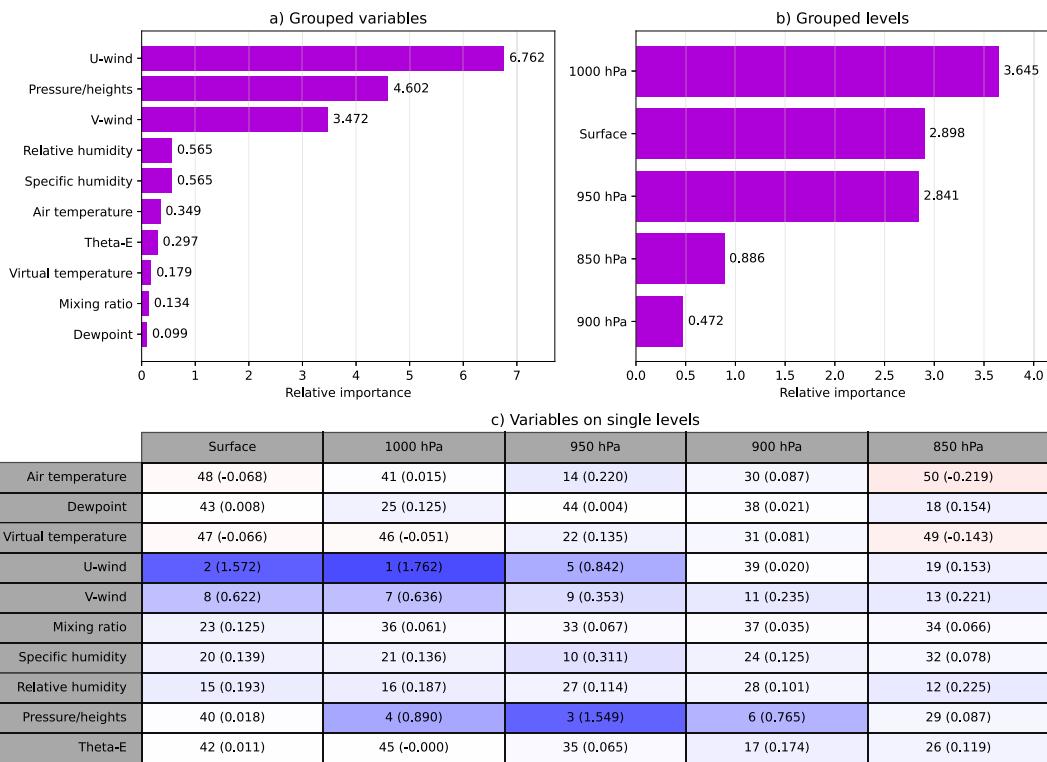


FIG. 11. Occluded front permutation results over the USAD for (a) grouped variables, (b) grouped vertical levels, and (c) variables on single levels ranked from 1 to 60 with 1 (60) being the most (least) important variable and level combination.

Occluded front performance was nearly the same across CONUS and USAD; 250-km CSI scores were about 0.54 over both domains. The spatial CSI plot in Figs. 10 and A4d shows a CSI maximum over the central United States east of the Rocky Mountains, with broader areas of high performance over much of the Pacific and Atlantic Ocean basins north of 30°N. The model appears to resolve the occlusion process for Norwegian model-like cyclones (Bjerknes and Solberg 1922), with occluded fronts wrapping around cyclones during the occlusion process (see the case study section for an example), in contrast to the Shapiro–Keyser cyclone model where an occluded front does not occur and the “wrap-up” process is associated with the bent-back warm front (Shapiro and Keyser 1990). However, the intersections of the cold, warm, and occluded fronts (commonly referred to as “triple points”) predicted by the model are not always collocated with the triple points shown in NOAA’s surface analyses. This could be due to the fact that most occlusions are warm-type occlusions, which occur when air behind former cold fronts is warmer than air ahead of the warm fronts (Schultz and Mass 1993; Schultz et al. 2014), and perhaps FrontFinder interprets portions of the warm-type occlusions as warm fronts.

The occluded front permutation studies in Figs. 11 and B4 show that wind and pressure variables were most important for occluded front detection, similar to our findings in J23

with the three-model system. This is supported by multiple studies showing wind shifts coincident with occluded fronts (e.g., Market and Moore 1998; Steenburgh et al. 1997; Schultz and Mass 1993; Shafer et al. 2006). In addition, the wind components are important in the wrap-up process, where an occluded front becomes elongated and wraps around the low pressure center cyclonically. This wrap-up is influenced by the deformation of the thermal ridge that is created after the cold and warm fronts collide (Schultz and Vaughan 2011; Martin 1999); however, it is unclear how much this contributed to the relative importance assigned to u wind and v wind as FrontFinder appears to struggle with identifying wrapped occlusions (discussed more in the case study section). Similar to cold fronts, occluded fronts are often associated with pressure troughs, and thus, sharp rises in pressure can be observed upon their passage (e.g., Shafer et al. 2006). The unique nature of occluded fronts being in close proximity to low pressure centers also likely contributes to greater relative importance for pressure variables. The most important vertical levels included the surface, 1000, and 950 hPa, which are the lowest vertical levels in our datasets and the first levels at which an occluded front will form following the collision of a cold front with a warm front.

Drylines over the Great Plains of the United States were detected well by FrontFinder. Using a 100-km neighborhood,

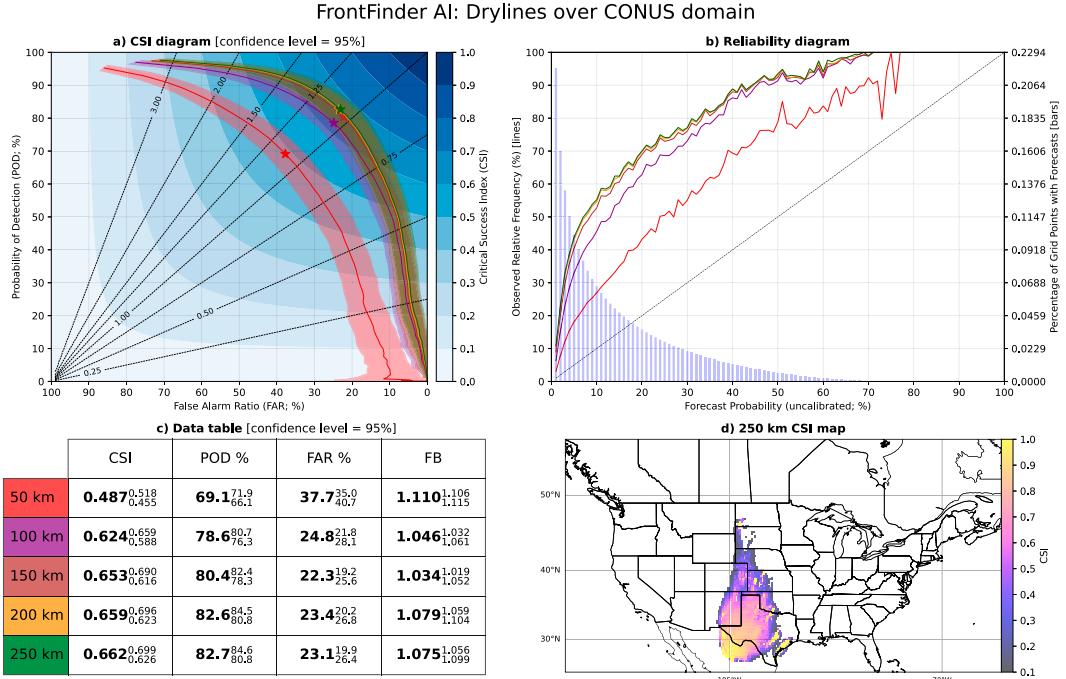


FIG. 12. Dryline results over CONUS: (a) CSI diagram (dashed lines = FB), (b) reliability diagram (dashed line = perfect reliability), (c) table with upper and lower performance bounds indicated with superscripts and subscripts, respectively, and (d) spatial CSI diagram using a 250-km neighborhood. CSI scores lower than 0.1 are not shown on the spatial diagram.

78.6% of drylines were located with a 24.8% FAR, yielding a 100-km CSI of 0.624 over the CONUS domain (Fig. 12). While we were encouraged to see the model perform well with drylines, which are known to have extreme moisture gradients at scales far smaller than the resolution of our training data, a UNET3+ model trained at higher resolutions may be able to reduce the FAR by narrowing its predictions along more well-defined drylines. While Clark et al. (2015) developed a dryline detection algorithm, the algorithm initially had a very high FAR (60%–70% with a 40-km neighborhood). However, their FAR dropped below 30% with the introduction of a random forest that assisted the algorithm, albeit with the trade-off of a 5%–10% decrease in POD. To our knowledge, FrontFinder is the first front detection algorithm that can predict and discern between drylines and cold, warm, stationary, and occluded fronts. Recent evidence has suggested that some of the false alarms from FrontFinder are due to changes in forecasters' drawing habits and that drylines were previously underdrawn by analysts (Hosek et al. 2025).

Our dryline permutation studies (Fig. 13) show that mixing ratio, v wind, and specific humidity were the most important variables for dryline detection. Drylines were the only front type predicted by the model to have a moisture variable as its most important predictor (mixing ratio), which is physically consistent with the fact that drylines are defined by their often extreme moisture gradients (Schaefer 1986). Specific humidity also had high importance, likely due to its similar conservation properties compared to those of the mixing ratio. The

high importance assigned to v wind was a surprising find; however, we believe this can be tied to the nocturnal low-level jet's southerly component that accelerates on the moist side of the dryline as the dryline begins its westward regression during the nocturnal hours (Parsons et al. 1991, 2000). Surface pressure also had elevated importance; observations by Parsons et al. (1991) and Ziegler and Hane (1993) showed that high (low) pressure exists on the dry (moist) side of the dryline. 900 and 850 hPa were determined to be the most important pressure levels for dryline detection, followed by the surface level. While we are unsure as to why higher pressure levels (in terms of elevation) were preferred over lower levels, we think this may be a result of diurnal heating effects that control the progression and regression of the dryline (Schaefer 1973) and the nocturnal low-level jet having a presence above the 900-hPa level as found by Parsons et al. (1991). While we think that a UNET3+ trained at higher resolutions would improve dryline performance, it is unclear how much of an effect (if any) this would have on results from our permutation studies.

#### 4. Case study: 2023 Christmas winter storm

With the goal of better understanding FrontFinder's strengths and weaknesses, we analyzed the model output (Fig. 14), middle-level water vapor satellite imagery (Fig. 15), WPC surface analyses (Fig. 16), and observations at the surface and 850 hPa (Figs. 17 and 18) for four time steps—0000 and 1200 UTC 26 and

## Dryline permutations: CONUS

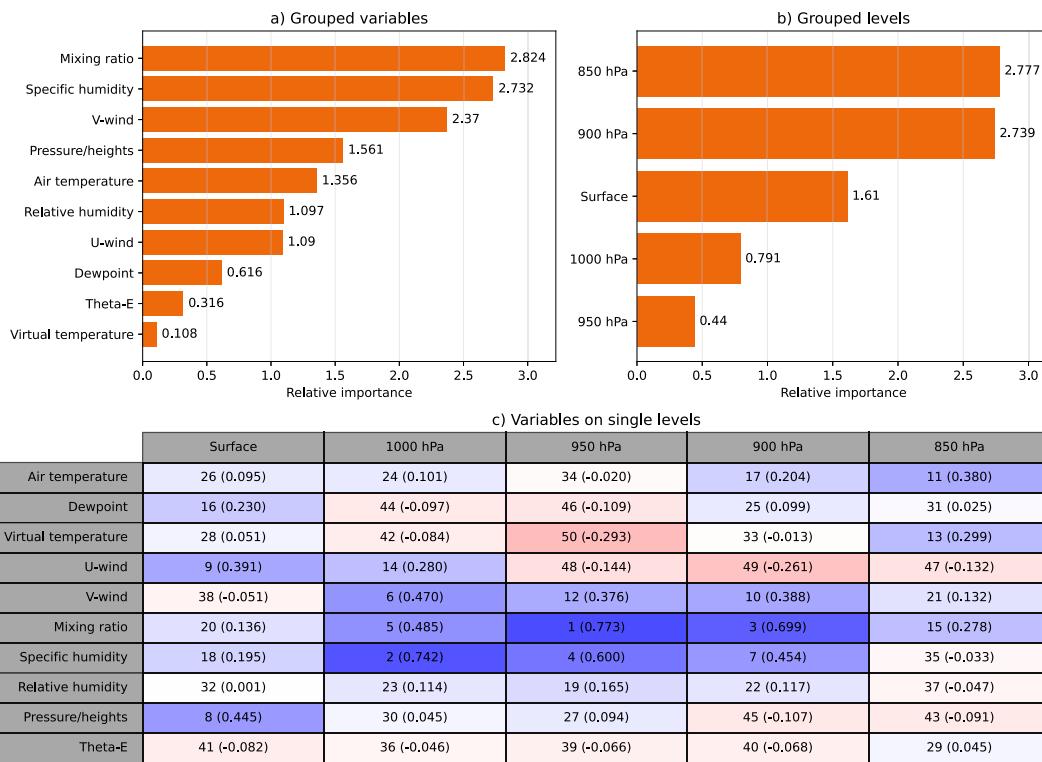


FIG. 13. As in Fig. B1, but for drylines.

27 December 2023. From 24 to 27 December 2023, a cyclone situated over the Midwest region of the United States brought upward of 12 in. (30.5 cm) of snow and 1.5 in. (3.8 cm) of ice across areas in South Dakota and Nebraska.<sup>1</sup> This precipitation was driven by an occluded front that was wrapped around its parent surface cyclone, with the precipitation occurring in the northwest quadrant of the cyclone. This is commonplace for winter precipitation as the northwest quadrants of surface cyclones are coincident with rising motion in frontogenetical circulations (Ganetis et al. 2018).

Starting at 0000 UTC 26 December 2023 (Figs. 14–18a), WPC surface analysis indicated an occluding cyclone over the upper Midwest with a nearby frontal zone extending from southern Minnesota to Quebec. The occluded front extends north of a triple point where a cold front is overtaking a warm front to the east as suggested by the WPC analysis and FrontFinder predictions. This occlusion was well defined with sharp temperature gradients at the surface and 850 hPa and a notable wind shift and pressure trough at the surface. FrontFinder had no issue reproducing the occluded front as calibrated probabilities exceeded 80% along most of the frontal zone (Fig. 14). The frontal zone extended from southern Minnesota to Quebec was analyzed by WPC as a stationary front in Minnesota and a cold front along the rest of its extent.

FrontFinder was able to highlight the transition from the stationary portion to the cold portion of the frontal zone, though the transition is displaced from that indicated in WPC's surface analysis. The model had lower probabilities along the stationary portion of the zone closest to the center of the cyclone. The model also identified a broad cold frontal region in the southeastern United States attached to the southern end of the occluded front. The model was likely less confident in this region due to a more diffuse temperature gradient, weaker wind shift, and the lack of a prominent pressure trough. In fact, this cold front has almost no presence at 850 hPa, though a weak moisture gradient can be noted where WPC analyzed a cold front. Satellite imagery shows the cyclone in the upper Midwest with an attendant high pressure system off the New England coast. Higher midlevel water vapor levels can be noted along and ahead of the occluded front, where precipitation is ongoing.

At 0300 UTC 26 December (not shown), WPC analyzed the previous occluded front as a cold front, likely due to air behind the front being much colder than air out ahead of it. Since this air was north of an analyzed warm front, this suggests that the occluded front analyzed 3 h prior may have been a rare instance of a cold-type occlusion, which occurs when air behind the former cold front undercuts warmer air ahead of the warm front (Schultz et al. 2014). Confirming this hypothesis would require an in-depth analysis of vertical cross sections along the occlusion, which is beyond the scope of this

<sup>1</sup> <https://www.weather.gov/abr/2023ChristmasStorm>.

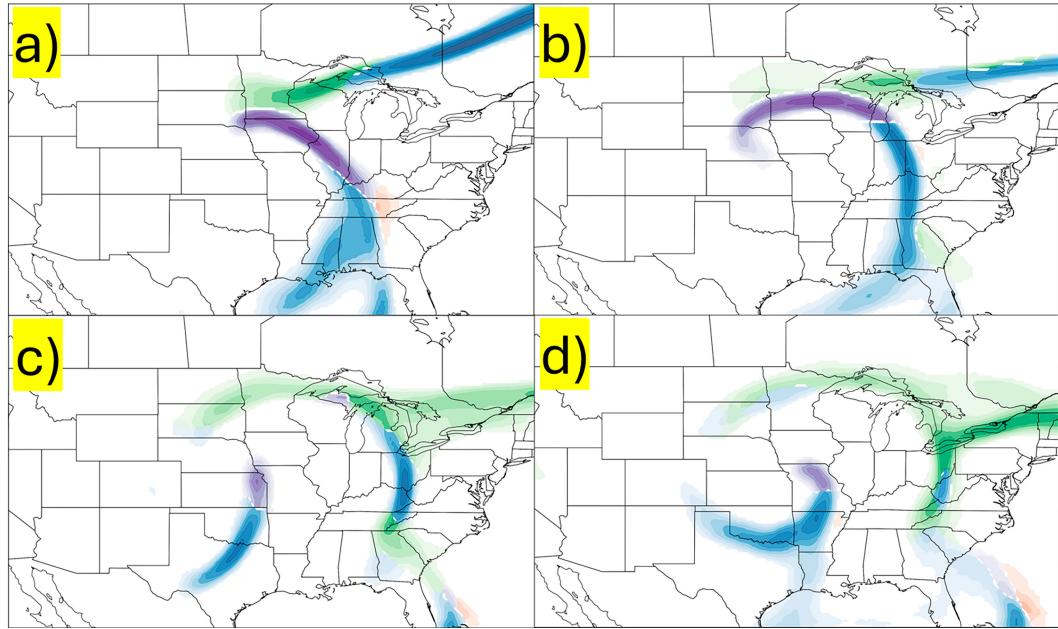


FIG. 14. FrontFinder predictions over CONUS using GFS model data: (a) 0000 UTC 26 Dec 2023, (b) 1200 UTC 26 Dec 2023, (c) 0000 UTC 27 Dec 2023, and (d) 1200 UTC 27 Dec 2023. Note that all predictions use forecast hour 0 for the respective initialization times and are calibrated to a 100-km neighborhood with filled contours at 10% intervals (blue = cold front, red = warm front, green = stationary front, purple = occluded front).

study. This front (now analyzed as a cold front) collided with a stationary boundary sitting in southern Minnesota, at which point a new occlusion was noted in the 0600 UTC 26 December WPC surface analysis (not shown). At 1200 UTC 26 December (Figs. 14–18b), WPC's surface analysis shows a triple point in western Wisconsin, with a stationary front extending to the

east-northeast (ENE), a cold front extending down to the Florida panhandle, and an occluded front that has started to wrap cyclonically around the low pressure center. The extension of the occluded front around the cyclone is likely due to deformation of the thermal ridge as described in [Schultz and Vaughan \(2011\)](#) and [Martin \(1999\)](#). The occluded front is highly

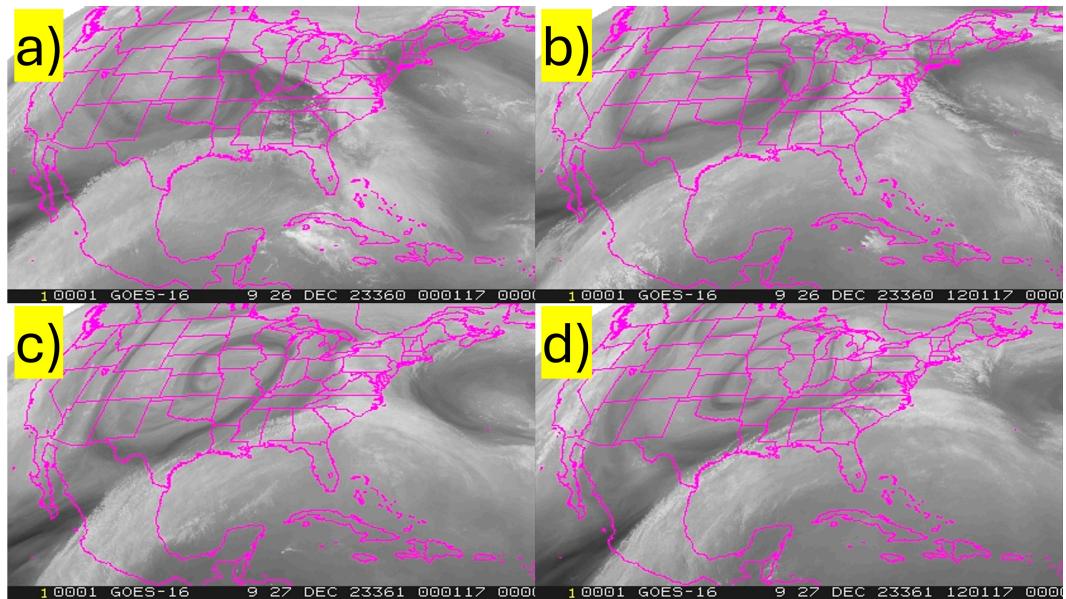


FIG. 15. GOES-16 midlevel water vapor imagery (band 9) showing a cyclone over CONUS: (a) 0001 UTC 26 Dec 2023, (b) 1201 UTC 26 Dec 2023, (c) 0001 UTC 27 Dec 2023, and (d) 1201 UTC 27 Dec 2023.

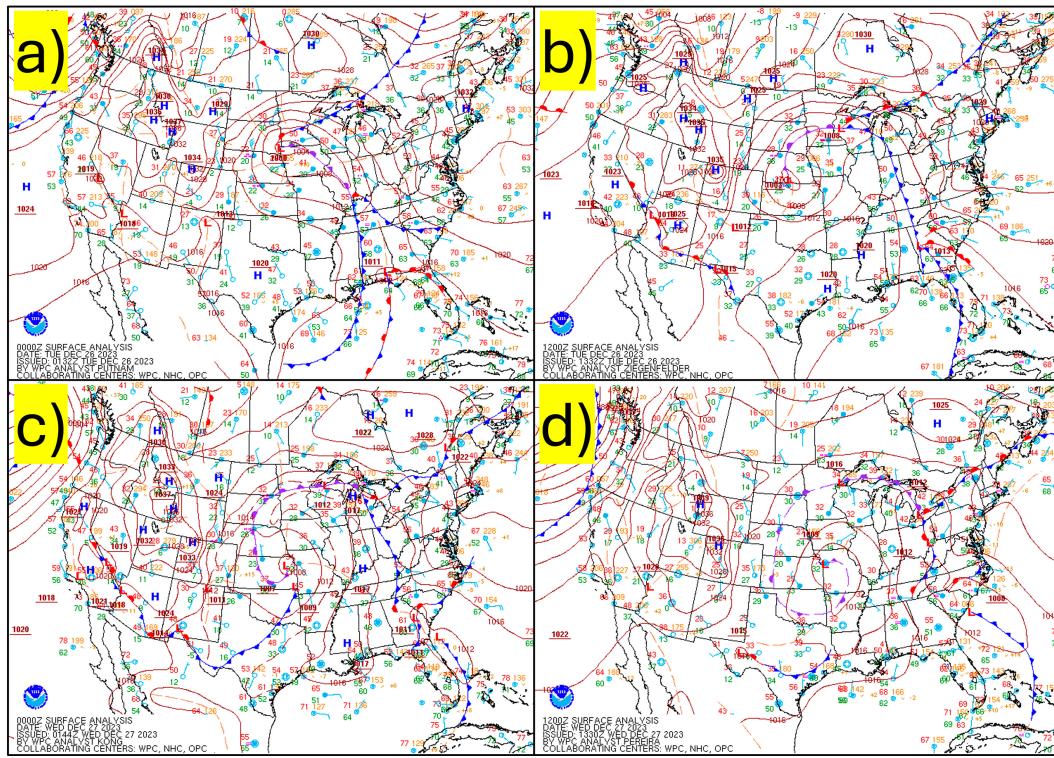


FIG. 16. WPC surface analyses over CONUS: (a) 0000 UTC 26 Dec 2023, (b) 1200 UTC 26 Dec 2023, (c) 0000 UTC 27 Dec 2023, and (d) 1200 UTC 27 Dec 2023.

visible on satellite in Minnesota and South Dakota with precipitation north and west of the front, highlighted by the midlevel water vapor imagery. FrontFinder was able to resolve the beginning of the wrap-up process, placing calibrated probabilities exceeding 70% around the northwest side of the low and identifying the occluded front responsible for heavy winter precipitation in the upper Midwest. Observations at the surface and 850 hPa show a wind shift and sharp temperature gradient along the occluded front, and a pressure trough is noted in the WPC surface analysis. The cold front across Ontario and Quebec had lower probabilities from the model, likely due to the front transitioning into a quasi-stationary state as suggested by the 1200 UTC surface analysis from WPC. This frontal zone was accompanied by a temperature gradient and strong convergence of v wind, consistent with our cold front permutation studies (see Figs. 4 and B1). The cold front extending down to the Florida panhandle also exhibited a temperature gradient and wind shift and was detected by FrontFinder with probabilities exceeding 80% in areas along the front. The model was less confident in its predictions closer to the Gulf of Mexico as a secondary surface low with multiple fronts had developed in the state of Georgia. The warm front analyzed by WPC extending east of this secondary low was not detected, even though we believe the surface observations corroborate WPC's analysis of the warm front. We believe that a combination of the weak temperature gradient along the warm front and the presence of other

boundaries in close proximity limited the model's ability to detect this warm front.

By 0000 UTC 27 December (Figs. 14–18c), the occluded front was wrapped three-quarters of the way around the cyclone, as evident by midlevel water vapor imagery. The dark corridor in the midlevel water vapor imagery behind the occluded front is representative of the dry conveyor belt, also known as the dry intrusion, which is a region of air that descends from the upper troposphere (Browning 1997). WPC analyzed this front along the transition between the dry conveyor belt and high midlevel water vapor concentrations coincident with winter precipitation. The occluded front had almost completely disappeared from the model output, with the exceptions of small corridors over the Upper Peninsula of Michigan and the Midwest, with the corridor in the Midwest being immediately south of the surface low and attached to a newly developed cold front in the southern Plains. Surface and 850-hPa analyses show that the temperature gradient along the occluded front has weakened significantly, though some wind convergence is still present along the occluded front in the Great Lakes region. FrontFinder had a mix of stationary, occluded, and cold front probabilities along the northernmost section of the analyzed occluded front, likely due to the wind convergence being accompanied by weak thermal gradients. The highest occluded front probabilities along this section of the analyzed occluded front were located near the triple point over Michigan, but FrontFinder

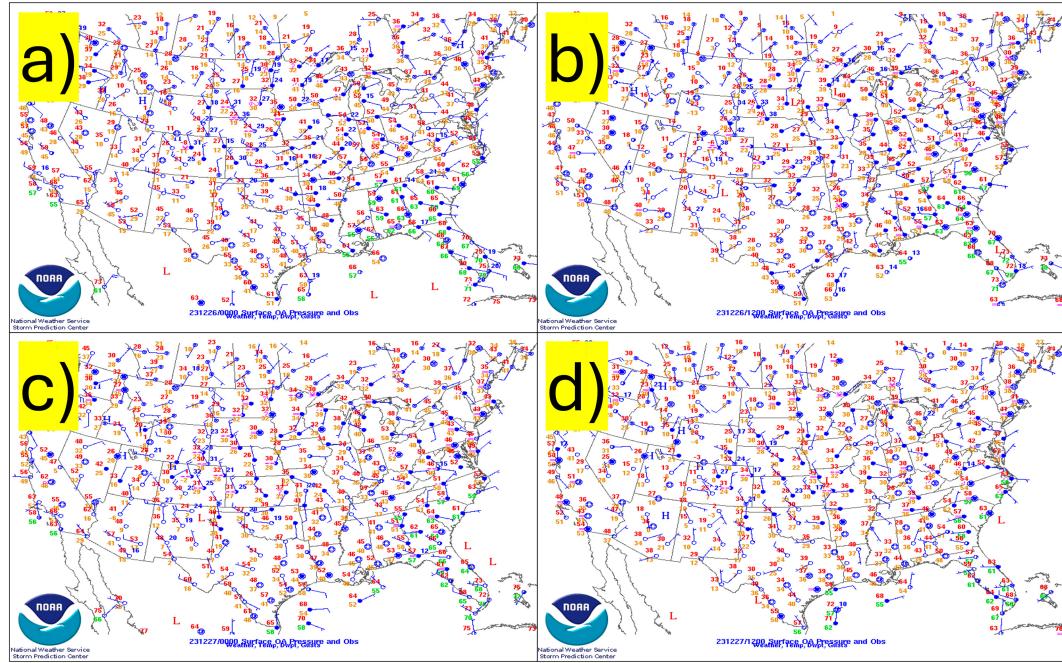


FIG. 17. Objectively analyzed surface maps from the Storm Prediction Center: (a) 0000 UTC 26 Dec 2023, (b) 1200 UTC 26 Dec 2023, (c) 0000 UTC 27 Dec 2023, and (d) 1200 UTC 27 Dec 2023.

still determined much of this frontal zone to be of the stationary type. At 1200 UTC 27 December (Figs. 14–18d), the occluded front is still visible on midlevel water vapor imagery and completely wrapped around the low pressure center. FrontFinder predictions did not exhibit any notable changes between 0000 and 1200 UTC 27 December; however, the occluded front probabilities near the triple point disappeared. While a subtle temperature gradient still exists at the surface and 850 hPa, there is no wind convergence that clearly demarcates the weakening occlusion. We believe that including satellite imagery as input to a future UNET3+ architecture will help with occluded front detection. Satellite imagery could also help with the detection of other front types as we learned that forecasters often use satellite imagery to locate fronts where surface observations are few and far between (e.g., oceans).

## 5. Conclusions and future work

We have shown that FrontFinder can effectively detect cold, warm, stationary, and occluded fronts and drylines over CONUS and NOAA's unified surface analysis domain. This model is a substantive improvement over our three-model system (TMS) highlighted in [J23](#), with greater fractions of fronts detected and lower false alarm ratios. FrontFinder outperforms TMS with cold, warm, stationary, and occluded fronts at the five neighborhoods assessed (50, 100, 150, 200, and 250 km), and all performance improvements were statistically significant at the 95% confidence level.

In terms of trainable parameters, FrontFinder is 25 times smaller than each of the models in TMS, which shows that

more parameters do not always result in better performance. Testing on five hardware devices (three GPUs and two CPUs) revealed that FrontFinder generates predictions 13–26 times faster than the TMS, with an average speed boost of 1693% across tested devices. It should be noted that this figure does not include data preprocessing; however, more details on the computational requirements of running FrontFinder in an operational environment and sample benchmarks can be found in [appendix C](#).

Combining all front types into a single model and the addition of drylines is likely a contributing factor to the performance improvements; FrontFinder is able to learn relationships and discern between all the front types, whereas TMS was unable to directly discern between certain front types because cold and warm fronts were separated from stationary and occluded fronts into their own models. Wet-bulb temperature and  $\theta_W$  were not included in the list of predictors for FrontFinder as permutation studies for TMS showed that these variables contributed very little to model performance; however, since FrontFinder is just one model, it is unclear if the exclusion of these variables had any impact on the model's performance. Assessing the true effect of wet-bulb temperature and  $\theta_W$  on FrontFinder's ability to detect fronts would require training of a separate UNET3+ architecture. Permutation studies performed on FrontFinder were consistent with those performed on TMS and showed that wind, temperature, and pressure and moisture variables contributed to FrontFinder's performance; however, some differences were noted, particularly with stationary fronts. Permutation studies performed on FrontFinder indicated that v wind negatively impacted stationary front performance, whereas TMS permutation studies from the model

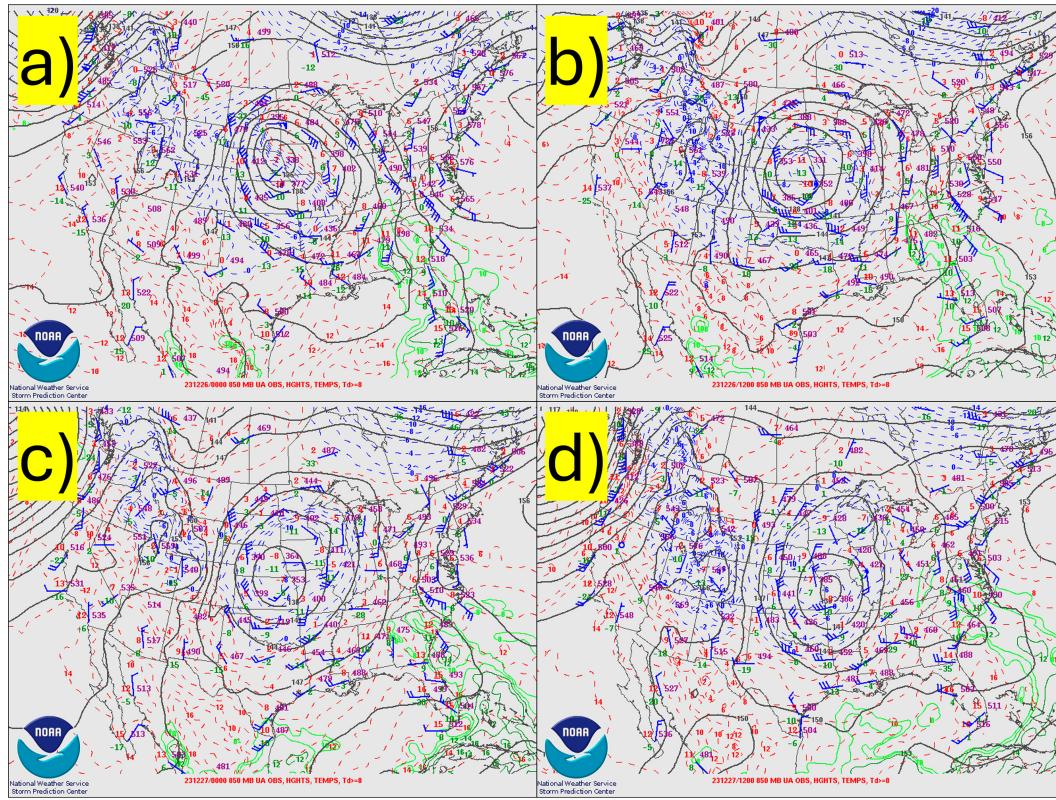


FIG. 18. Objectively analyzed 850-hPa maps from the Storm Prediction Center: (a) 0000 UTC 26 Dec 2023, (b) 1200 UTC 26 Dec 2023, (c) 0000 UTC 27 Dec 2023, and (d) 1200 UTC 27 Dec 2023.

predicting stationary and occluded fronts suggested that  $v$  wind was an important predictor for stationary fronts. Discrepancies between the results of the permutation studies from FrontFinder and the TMS could be explained by the “peaking phenomenon” described in Sima and Dougherty (2008). Since we decreased the number of variables from the TMS to FrontFinder by removing the wet-bulb temperature and wet-bulb potential temperature, this could have resulted in poorer classification accuracy compared to a model that includes these variables as inputs. We plan to improve upon our permutation studies by implementing temporal and feature-based cross-validation strategies in future work (e.g., Marco et al. 2022; Roberts et al. 2017). Nonetheless, we believe that results from permutation studies conducted on FrontFinder are more representative of variables’ true contributions to front detection since FrontFinder is able to directly learn relationships between the five front types it was trained to predict.

We also think that accounting for correlations between the input variables through greater data augmentation could improve the performance of future model architectures. The case study of the 2023 Christmas Winter Storm showed that FrontFinder can resolve the early stages of an occluding extratropical cyclone; however, satellite data might also be needed to train a model that can fully resolve the wrap-up of some occluded fronts.

FrontFinder is now used operationally at WPC, and over the coming months, we hope to get constructive feedback on how effectively forecasters are able to use our new algorithm to improve the identification and classification of frontal boundaries. This feedback will help us make important decisions about how we structure and train future algorithms for identifying fronts. We are also collaborating with forecasters at The Weather Company (TWC) to assist them in drawing fronts globally out to 120-h lead time, a labor-intensive process that motivated TWC to adopt our algorithm and use real-time Global Forecast System (GFS) and European Centre for Medium-Range Weather Forecasts (ECMWF) model data as input to FrontFinder to highlight frontal zones across the globe. While no performance statistics have been calculated globally due to the lack of available ground truth front labels, initial impressions of performance in the Southern Hemisphere are poor, suggesting that we need a larger domain of training data in order to resolve fronts outside of USAD, consistent with other studies that have suggested similar issues with regional training (e.g., Niebler et al. 2022). We hope to eventually develop an algorithm that is able to resolve fronts across the globe with spatially uniform performance; such an algorithm could enable us to perform long-term climatologies of different types of frontal boundaries across the globe and integrate this model operationally at many different international weather information providers.

Our future work will include training new model architectures on high-resolution model data such as the North American

Mesoscale Forecast System (NAM), High-Resolution Rapid Refresh (HRRR) model, and the Rapid Refresh Forecast System (RRFS) model that is currently under development. HRRR and RRFS are convective-allowing models (CAMs) on 3-km grids, so small-scale convective feedbacks will likely complicate the training of future algorithms with HRRR and RRFS data as input. However, a front detection algorithm trained on high-resolution data could have benefits over ERA5 reanalysis, such as being able to precisely locate boundaries that can influence storm-scale processes like tornadogenesis (e.g., Maddox et al. 1980; Markowski et al. 1998; Sills et al. 2004; Childs and Schumacher 2019). With the exponential growth of ML in meteorological research (Chase et al. 2022), future high-resolution deep learning models could also be used to train other model architectures to predict fronts to ascertain the performance of these important features. Many recently developed ML algorithms have been able to outperform traditional numerical weather prediction at forecasting precipitation and variables such as temperature, specific humidity, wind speed, and geopotential; however, whether processes such as fronts are being appropriately resolved remains an open question (e.g., Sønderby et al. 2020; Bi et al. 2022; Andrychowicz et al. 2023; Zhang et al. 2023). Permutation studies like those performed in this paper could be used to increase our understanding of fronts and improve existing NWP models, eventually providing the basis for better reanalysis datasets that can be used to train future machine learning models.

**Acknowledgments.** This material is based on work supported by the National Science Foundation under Grant RISE-2019758. This material is also based on work supported by the National Oceanic and Atmospheric Administration under Grant NA20OAR4590347. We thank WPC, OPC, TAFB, and TWC forecasters and analysts for providing feedback on our algorithm and working with us to operationalize the frontal model. The results contain modified Copernicus Climate Change Service Information 2021. Neither the European Commission nor ECMWF is responsible for any use that may be made of the Copernicus information or data it contains.

**Data availability statement.** ERA5 data on single and pressure levels were downloaded from the Copernicus Climate Change Service (C3S) Climate Data Store and can be found via Hersbach et al. (2023). Frontal data derived from the WPC analyses can be found via NOAA (2023). Python code used in this project is available in the “AIES-D-24-0043” branch of our GitHub repository at <https://github.com/ai2es/fronts/tree/AIES-D-24-0043>.

## APPENDIX A

### Performance Diagrams: CONUS

Figures A1–A4 show FrontFinder’s performance diagrams for cold, warm, stationary, and occluded fronts over our CONUS domain.

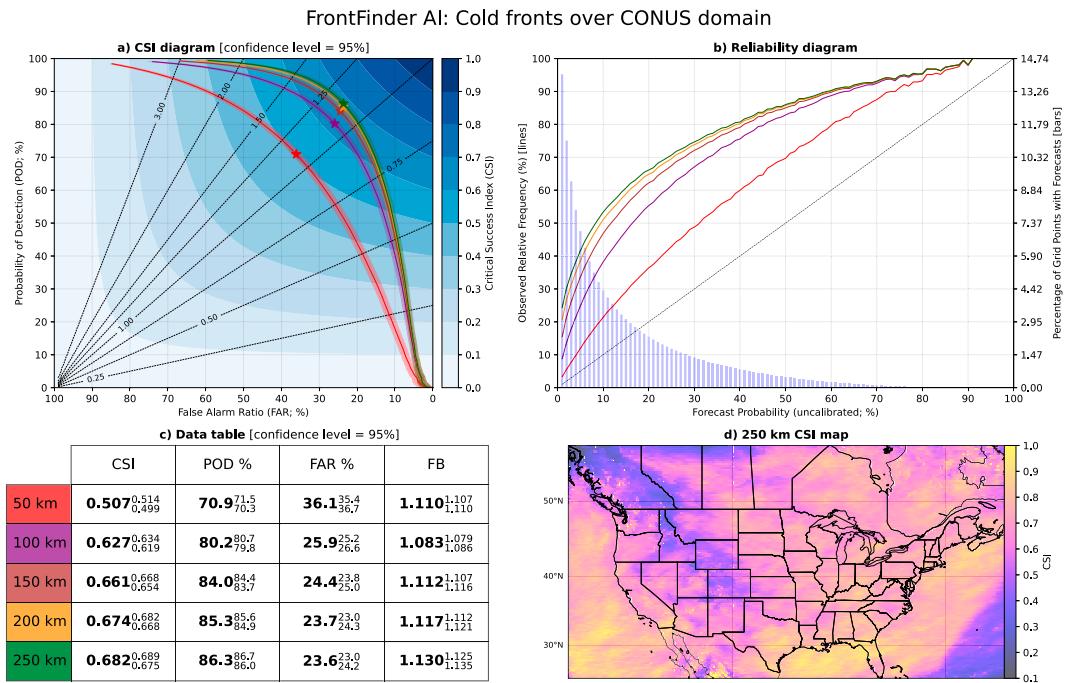


FIG. A1. Cold front results over CONUS: (a) CSI diagram (dashed lines = FB), (b) reliability diagram (dashed line = perfect reliability), (c) table with upper and lower performance bounds indicated with superscripts and subscripts, respectively, and (d) spatial CSI diagram using a 250-km neighborhood.

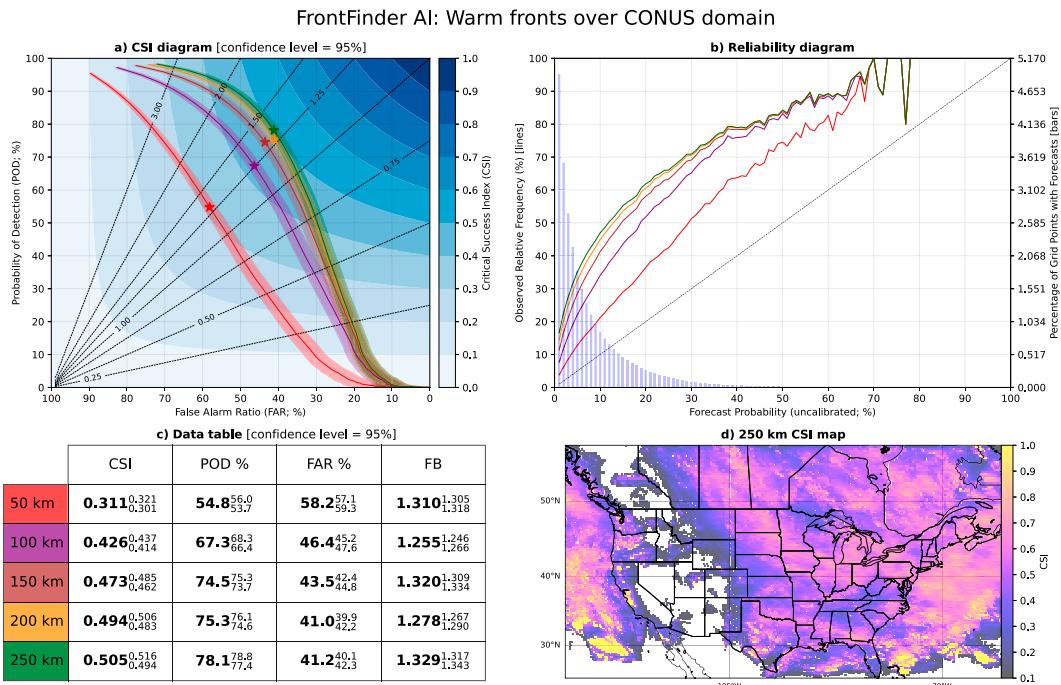


FIG. A2. Warm front results over CONUS: (a) CSI diagram (dashed lines = FB), (b) reliability diagram (dashed line = perfect reliability), (c) table with upper and lower performance bounds indicated with superscripts and subscripts, respectively, and (d) spatial CSI diagram using a 250-km neighborhood.

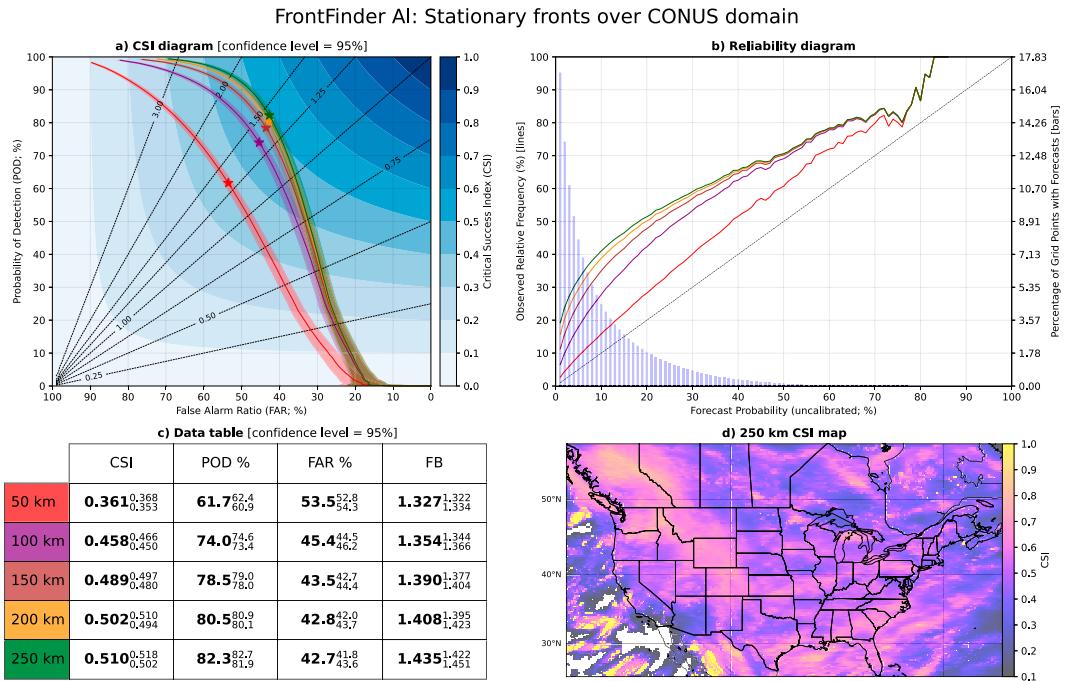


FIG. A3. Stationary front results over CONUS: (a) CSI diagram (dashed lines = FB), (b) reliability diagram (dashed line = perfect reliability), (c) table with upper and lower performance bounds indicated with superscripts and subscripts, respectively, and (d) spatial CSI diagram using a 250-km neighborhood.

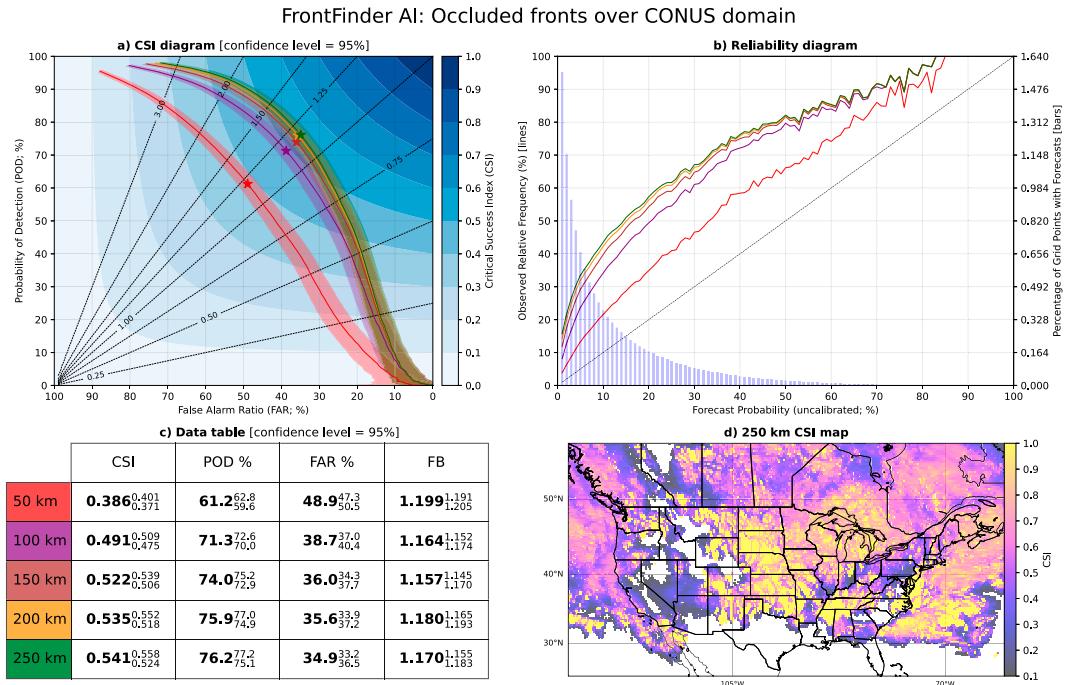


FIG. A4. Occluded front results over CONUS: (a) CSI diagram (dashed lines = FB), (b) reliability diagram (dashed line = perfect reliability), (c) table with upper and lower performance bounds indicated with superscripts and subscripts, respectively, and (d) spatial CSI diagram using a 250-km neighborhood.

## APPENDIX B

### Permutation Studies: CONUS

Figures B1–B4 show permutation studies performed on FrontFinder for cold, warm, stationary, and occluded fronts over our CONUS domain.

## Cold front permutations: CONUS

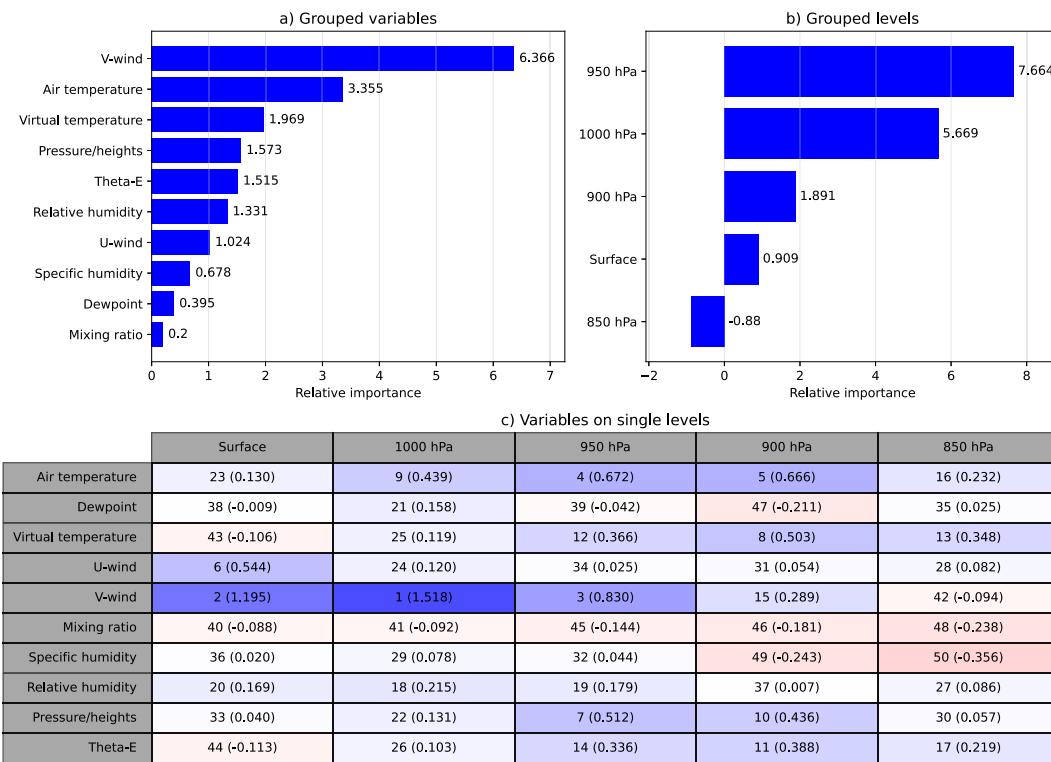


FIG. B1. Cold front permutation results over the CONUS domain for (a) grouped variables, (b) grouped vertical levels, and (c) variables on single levels ranked from 1 to 60 with 1 (60) being the most (least) important variable and level combination.

## Warm front permutations: CONUS

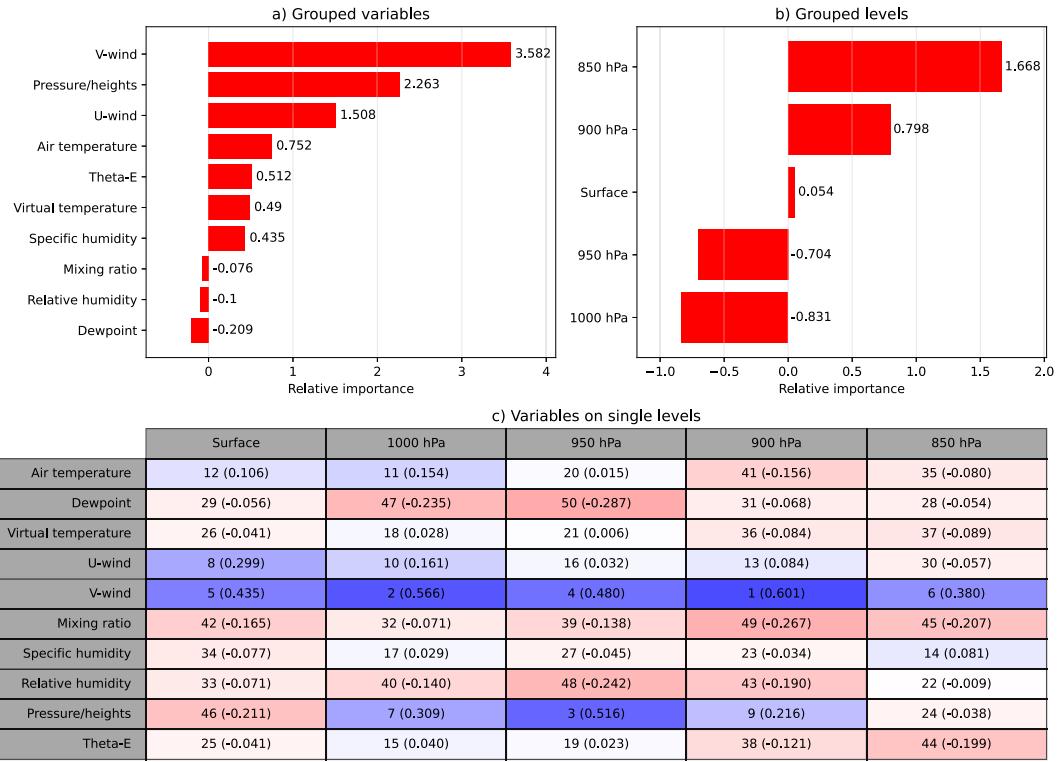


FIG. B2. Warm front permutation results over the CONUS domain for (a) grouped variables, (b) grouped vertical levels, and (c) variables on single levels ranked from 1 to 60 with 1 (60) being the most (least) important variable and level combination.

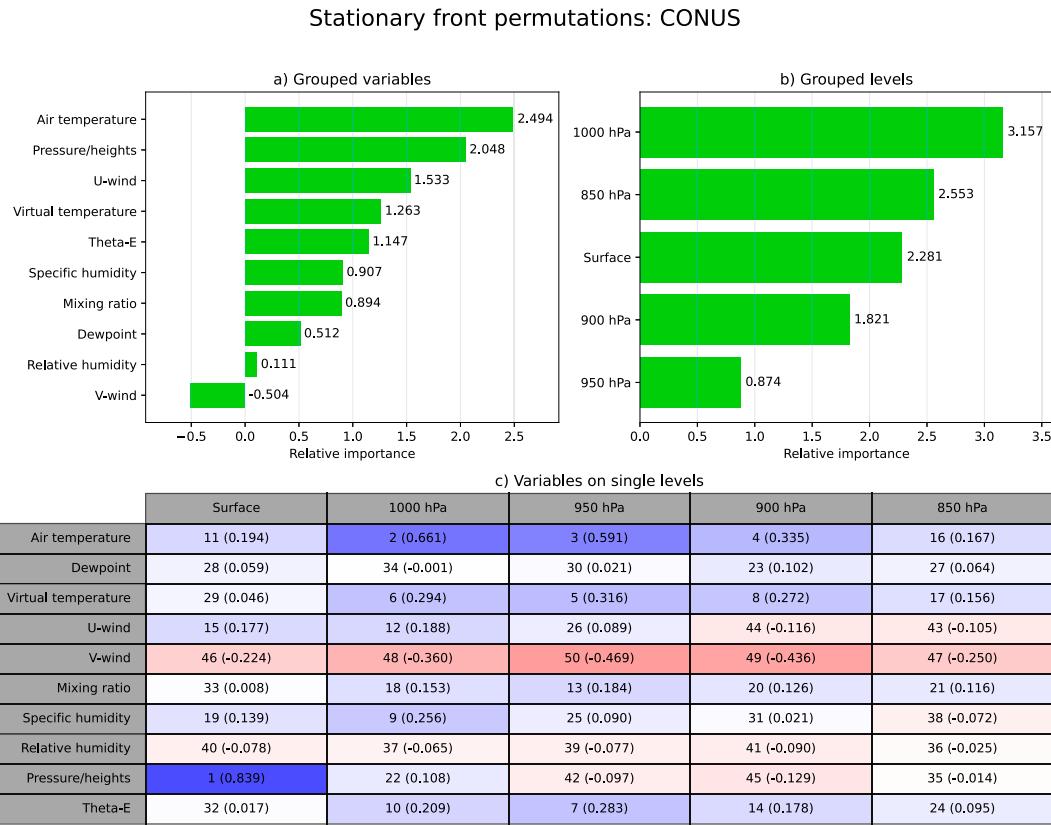


FIG. B3. Stationary front permutation results over the CONUS domain for (a) grouped variables, (b) grouped vertical levels, and (c) variables on single levels ranked from 1 to 60 with 1 (60) being the most (least) important variable and level combination.

## Occluded front permutations: CONUS

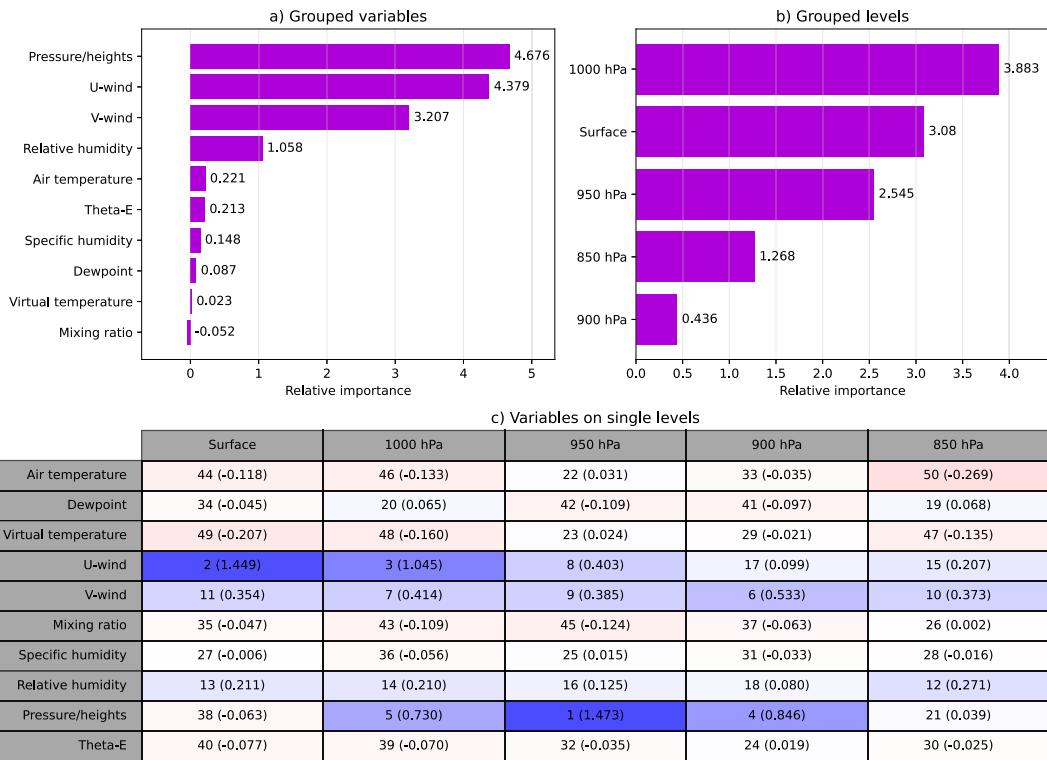


FIG. B4. Occluded front permutation results over the CONUS domain for (a) grouped variables, (b) grouped vertical levels, and (c) variables on single levels ranked from 1 to 60 with 1 (60) being the most (least) important variable and level combination.

## APPENDIX C

## Computational Requirements and Benchmarks

In this section, we analyze the computational overhead associated with running FrontFinder in operational environments. For the purposes of this study, we calculated benchmarks in an environment that aims to make predictions with FrontFinder over USAD for eight forecast time steps during a single run. Comparisons are made between the workflows of running FrontFinder and the TMS operationally. A summary of expected computation times required for the FrontFinder workflow can be found in Table C1 at the end of this section. Note that quoted memory and time requirements should be treated as estimates; these can vary significantly between systems.

The inputs to FrontFinder are sourced from the GFS model, the data for which are updated in real time and publicly available via an Amazon S3 bucket. Downloading eight forecast time steps for a single initialization time will require about 4.2 GB of local hard disk space (just over 0.5 GB per time step). Network speed is the primary bottleneck in running FrontFinder operationally as a slow network can cause delays in the retrieval of GFS data.

Data preprocessing can be broken up into two parts: reading the data and calculating additional variables. The speed at

which the retrieved GFS datasets can be opened depends on many factors including, but not limited to, the utilized storage device, CPU, and RAM. Testing across three different systems suggests that reading all retrieved data can take anywhere from 18 s to 4 min, with a median time of 2 min and 28 s. Calculating additional variables after the data has been read only takes 4–7 s for both FrontFinder and the TMS. Data preprocessing requires the most RAM of any stage of the workflow; users should expect to use up to 6 GB of RAM throughout the entire workflow. The main bottleneck of the data preprocessing

TABLE C1. Estimated computation time for stages of the FrontFinder workflow in an operational environment aiming to generate predictions for eight forecast time steps over USAD. Note that the minimum, median, and maximum times for data retrieval assume network download speeds of 100, 50, and 10 MB s<sup>-1</sup>, respectively (m = minutes; s = seconds). Values in boldface represent the total processing time.

Workflow stage	Minimum time	Median time	Maximum time
Data retrieval	43 s	1 m 26 s	7 m 10 s
Data preprocessing	22 s	2 m 34 s	4 m 7 s
Predictions (GPU)	10 s	21 s	29 s
Predictions (CPU)	49 s	1 m 22 s	1 m 55 s
<b>Total (GPU)</b>	<b>1 m 11 s</b>	<b>4 m 19 s</b>	<b>11 m 46 s</b>
<b>Total (CPU)</b>	<b>1 m 50 s</b>	<b>5 m 20 s</b>	<b>13 m 12 s</b>

stage of the workflow is the speed at which the GFS datasets can be opened.

Generating predictions with FrontFinder and the TMS can be accomplished on either a GPU or a CPU. Testing on three different GPUs (NVIDIA Tesla A100, NVIDIA RTX 4070 SUPER, NVIDIA TITAN XP) suggests that GPU users can expect a set of eight FrontFinder predictions over USAD to take 10–29 s, with a median prediction time of 21 s. In contrast, the same set of predictions with the TMS on a GPU requires between 80 s and 6 min of computation time, with a median time of 4 min and 27 s across tested GPU devices. Testing on two different CPUs (Intel Xeon E5-2670 v3 and AMD Ryzen 9 5900X, using 12 cores on each) suggests that CPU users can expect FrontFinder to spend 49–115 s generating a set of eight predictions over USAD, with a median time of 82 s across tested CPUs. The same set of predictions with the TMS on a CPU takes 20–31 min (median of 25.5 min), roughly 19 times slower than FrontFinder.

## REFERENCES

Andrychowicz, M., L. Espeholt, D. Li, S. Merchant, A. Merose, F. Zyda, S. Agrawal, and N. Kalchbrenner, 2023: Deep learning for day forecasts from sparse observations. arXiv, 2306.06079v3, <https://doi.org/10.48550/arXiv.2306.06079>.

Berry, G., M. J. Reeder, and C. Jakob, 2011: A global climatology of atmospheric fronts. *Geophys. Res. Lett.*, **38**, L04809, <https://doi.org/10.1029/2010GL046451>.

Bi, K., L. Xie, H. Zhang, X. Chen, X. Gu, and Q. Tian, 2022: Pangu-Weather: A 3D high-resolution model for fast and accurate global weather forecast. arXiv, 2211.02556v1, <https://doi.org/10.48550/arXiv.2211.02556>.

Biard, J. C., and K. E. Kunkel, 2019: Automated detection of weather fronts using a deep learning neural network. *Adv. Stat. Climatol. Meteor. Oceanogr.*, **5**, 147–160, <https://doi.org/10.5194/ascmo-5-147-2019>.

Bjerknes, J., 1919: On the structure of moving cyclones. *Mon. Wea. Rev.*, **47**, 95–99, [https://doi.org/10.1175/1520-0493\(1919\)47<95:OTSOMC>2.0.CO;2](https://doi.org/10.1175/1520-0493(1919)47<95:OTSOMC>2.0.CO;2).

—, and H. Solberg, 1922: Life cycle of cyclones and the polar front theory of atmospheric circulation. *Geofys. Publ.*, **12**, 1–61.

Bochenek, B., Z. Ustrnul, A. Wypych, and D. Kubacka, 2021: Machine learning-based front detection in Central Europe. *Atmosphere*, **12**, 1312, <https://doi.org/10.3390/atmos12101312>.

Breiman, L., 2001: Random forests. *Mach. Learn.*, **45**, 5–32, <https://doi.org/10.1023/A:1010933404324>.

Bridle, J. S., 1989: Training stochastic model recognition algorithms as networks can lead to maximum mutual information estimation of parameters. *NIPS'89: Proceedings of the 3rd International Conference on Neural Information Processing Systems*, MIT Press, 211–217, <https://dl.acm.org/doi/10.5555/2969830.2969856>.

Browning, K. A., 1997: The dry intrusion perspective of extra-tropical cyclone development. *Meteor. Appl.*, **4**, 317–324, <https://doi.org/10.1017/S1350482797000613>.

Chase, R. J., D. R. Harrison, A. Burke, G. M. Lackmann, and A. McGovern, 2022: A machine learning tutorial for operational meteorology. Part I: Traditional machine learning. *Wea. Forecasting*, **37**, 1509–1529, <https://doi.org/10.1175/WAF-D-22-0070.1>.

Childs, S. J., and R. S. Schumacher, 2019: An updated severe hail and tornado climatology for eastern Colorado. *J. Appl. Meteor. Climatol.*, **58**, 2273–2293, <https://doi.org/10.1175/JAMC-D-19-0098.1>.

Clark, A. J., A. MacKenzie, A. McGovern, V. Lakshmanan, and R. A. Brown, 2015: An automated, multiparameter dryline identification algorithm. *Wea. Forecasting*, **30**, 1781–1794, <https://doi.org/10.1175/WAF-D-15-0070.1>.

Clarke, L. C., and R. J. 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).

Denker, J. S., and C. C. J. Burges, 1995: Image segmentation and recognition. *The Mathematics of Generalization*, CRC Press, 409–434.

Ganetis, S. A., B. A. Colle, S. E. Yuter, and N. P. Hoban, 2018: Environmental conditions associated with observed snowband structures within northeast U.S. winter storms. *Mon. Wea. Rev.*, **146**, 3675–3690, <https://doi.org/10.1175/MWR-D-18-0054.1>.

Hendrycks, D., and K. Gimpel, 2023: Gaussian Error Linear Units (GELUs). arXiv, 1606.08415v5, <https://doi.org/10.48550/arXiv.1606.08415>.

Hersbach, H., and Coauthors, 2023: ERA5 hourly data on pressure levels from 1940 to present. Copernicus Climate Change Service (C3S) Climate Data Store (CDS), accessed 15 March 2021, <https://doi.org/10.24381/cds.bd0915c6>.

Hewson, T. D., 1998: Objective fronts. *Meteor. Appl.*, **5**, 37–65, <https://doi.org/10.1017/S1350482798000553>.

Hines, K. M., and C. R. Mechoso, 1993: Influence of surface drag on the evolution of fronts. *Mon. Wea. Rev.*, **121**, 1152–1176, [https://doi.org/10.1175/1520-0493\(1993\)121<1152:IOSDOT>2.0.CO;2](https://doi.org/10.1175/1520-0493(1993)121<1152:IOSDOT>2.0.CO;2).

Hosek, M. J., K. A. Hoogewind, A. J. Clark, A. D. Justin, and J. T. Allen, 2025: A 16-year climatology of WPC-analyzed drylines and their association with severe convection. *J. Appl. Meteor. Climatol.*, in press.

Huang, H., and Coauthors, 2020: UNet 3+: A full-scale connected UNet for medical image segmentation. arXiv, 2004.08790v1, <https://doi.org/10.48550/arxiv.2004.08790>.

Ioffe, S., and C. Szegedy, 2015: Batch normalization: Accelerating deep network training by reducing internal covariate shift. arXiv, 1502.03167v3, <https://doi.org/10.48550/arXiv.1502.03167>.

Justin, A., 2024: Explainable frontal boundary predictions for applications in operational environments. M.S. thesis, Dept. of Meteorology, University of Oklahoma, 105 pp., <https://hdl.handle.net/11244/340396>.

Justin, A. D., C. Willingham, A. McGovern, and J. T. Allen, 2023: Toward operational real-time identification of frontal boundaries using machine learning. *Artif. Intell. Earth Syst.*, **2**, e220052, <https://doi.org/10.1175/AIES-D-22-0052.1>.

Kingma, D. P., and J. Ba, 2014: Adam: A method for stochastic optimization. arXiv, 1412.6980v9, <https://doi.org/10.48550/arxiv.1412.6980>.

Kunkel, K. E., D. R. Easterling, D. A. R. Kristovich, B. Gleason, L. Stoecker, and R. Smith, 2012: Meteorological causes of the secular variations in observed extreme precipitation events for the conterminous United States. *J. Hydrometeor.*, **13**, 1131–1141, <https://doi.org/10.1175/JHM-D-11-0108.1>.

Lagerquist, R., A. McGovern, and D. J. 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>.

—, J. T. Allen, and A. McGovern, 2020: Climatology and variability of warm and cold fronts over North America from 1979 to 2018. *J. Climate*, **33**, 6531–6554, <https://doi.org/10.1175/JCLI-D-19-0680.1>.

LeCun, Y., B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel, 1989: Backpropagation applied to handwritten zip code recognition. *Neural Comput.*, **1**, 541–551, <https://doi.org/10.1162/neco.1989.1.4.541>.

Lee, M., 2023: GELU activation function in deep learning: A comprehensive mathematical analysis and performance. arXiv, 2305.12073v2, <https://doi.org/10.48550/arXiv.2305.12073>.

Lu, L., Y. Shin, Y. Su, and G. E. Karniadakis, 2020: Dying ReLU and initialization: Theory and numerical examples. *Commun. Comput. Phys.*, **28**, 1671–1706, <https://doi.org/10.4208/cicp.OA-2020-0165>.

Maddox, R. A., L. R. Hoxit, and C. F. Chappell, 1980: A study of tornadic thunderstorm interactions with thermal boundaries. *Mon. Wea. Rev.*, **108**, 322–336, [https://doi.org/10.1175/1520-0493\(1980\)108<0322:ASOTTI>2.0.CO;2](https://doi.org/10.1175/1520-0493(1980)108<0322:ASOTTI>2.0.CO;2).

Marco, Z., A. Elena, S. Anna, T. Silvia, and C. Andrea, 2022: Spatio-temporal cross-validation to predict pluvial flood events in the Metropolitan City of Venice. *J. Hydrol.*, **612**, 128150, <https://doi.org/10.1016/j.jhydrol.2022.128150>.

Market, P. S., and J. T. Moore, 1998: Mesoscale evolution of a continental occluded cyclone. *Mon. Wea. Rev.*, **126**, 1793–1811, [https://doi.org/10.1175/1520-0493\(1998\)126<1793:MEOACO>2.0.CO;2](https://doi.org/10.1175/1520-0493(1998)126<1793:MEOACO>2.0.CO;2).

Markowski, P. M., E. N. Rasmussen, and J. M. Straka, 1998: The occurrence of tornadoes in supercells interacting with boundaries during VORTEX-95. *Wea. Forecasting*, **13**, 852–859, [https://doi.org/10.1175/1520-0434\(1998\)013<0852:TOOTIS>2.0.CO;2](https://doi.org/10.1175/1520-0434(1998)013<0852:TOOTIS>2.0.CO;2).

Martin, J. E., 1999: The separate roles of geostrophic vorticity and deformation in the midlatitude occlusion process. *Mon. Wea. Rev.*, **127**, 2404–2418, [https://doi.org/10.1175/1520-0493\(1999\)127<2404:TSROGV>2.0.CO;2](https://doi.org/10.1175/1520-0493(1999)127<2404:TSROGV>2.0.CO;2).

Matsuoka, D., S. Sugimoto, Y. Nakagawa, S. Kawahara, F. Araki, Y. Onoue, M. Iiyama, and K. Koyamada, 2019: Automatic detection of stationary fronts around Japan using a deep convolutional neural network. *SOLA*, **15**, 154–159, <https://doi.org/10.2151/sola.2019-028>.

McGovern, A., R. Lagerquist, D. J. Gagne II, G. E. Jergensen, K. L. Elmore, C. R. 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>.

Mitchell, K. E., and J. B. Hovermale, 1977: A numerical investigation of the severe thunderstorm gust front. *Mon. Wea. Rev.*, **105**, 657–675, [https://doi.org/10.1175/1520-0493\(1977\)105<0657:ANIOTS>2.0.CO;2](https://doi.org/10.1175/1520-0493(1977)105<0657:ANIOTS>2.0.CO;2).

National Weather Service, 2019: National Weather Service coded surface bulletins, 2003-. Zenodo, accessed 15 March 2021, <https://doi.org/10.5281/zenodo.2642801>.

Niebler, S., A. Miltenberger, B. Schmidt, and P. Spichtinger, 2022: Automated detection and classification of synoptic-scale fronts from atmospheric data grids. *Wea. Climate Dyn.*, **3**, 113–137, <https://doi.org/10.5194/wcd-3-113-2022>.

NOAA, 2023: NOAA unified surface analysis fronts. Zenodo, accessed 5 January 2023, <https://doi.org/10.5281/zenodo.7505022>.

O’Shea, K., and R. Nash, 2015: An introduction to convolutional neural networks. arXiv, 1511.08458v2, <https://doi.org/10.48550/arxiv.1511.08458>.

Parsons, D. B., M. A. Shapiro, R. M. Hardesty, R. J. Zamora, and J. M. Intrieri, 1991: The finescale structure of a West Texas dryline. *Mon. Wea. Rev.*, **119**, 1242–1258, [https://doi.org/10.1175/1520-0493\(1991\)119<1242:TFSOAW>2.0.CO;2](https://doi.org/10.1175/1520-0493(1991)119<1242:TFSOAW>2.0.CO;2).

—, —, and E. Miller, 2000: The mesoscale structure of a nocturnal dryline and of a frontal–dryline merger. *Mon. Wea. Rev.*, **128**, 3824–3838, [https://doi.org/10.1175/1520-0493\(2001\)129<3824:TMSOAN>2.0.CO;2](https://doi.org/10.1175/1520-0493(2001)129<3824:TMSOAN>2.0.CO;2).

Pietrycha, A. E., and E. N. Rasmussen, 2004: Finescale surface observations of the dryline: A mobile mesonet perspective. *Wea. Forecasting*, **19**, 1075–1088, <https://doi.org/10.1175/819.1>.

Renard, R. J., and L. C. Clarke, 1965: Experiments in numerical objective frontal analysis. *Mon. Wea. Rev.*, **93**, 547–556, [https://doi.org/10.1175/1520-0493\(1965\)093%3C0547:EINOFA%3E2.3.CO;2](https://doi.org/10.1175/1520-0493(1965)093%3C0547:EINOFA%3E2.3.CO;2).

Roberts, D. R., and Coauthors, 2017: Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. *Ecography*, **40**, 913–929, <https://doi.org/10.1111/ecog.02881>.

Roberts, N., 2008: Assessing the spatial and temporal variation in the skill of precipitation forecasts from an NWP model. *Met. Appl.*, **15**, 163–169, <https://doi.org/10.1002/met.57>.

Ronneberger, O., P. Fischer, and T. Brox, 2015: U-Net: Convolutional networks for biomedical image segmentation. *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015*, N. Navab et al., Eds., Springer, 234–241.

Schaefer, J. T., 1973: The motion and morphology of the dryline. NOAA Tech. Memo. ERL NSSL-66, 81 pp., <https://repository.library.noaa.gov/view/noaa/19283>.

—, 1986: The dryline. *Mesoscale Meteorology and Forecasting*, Amer. Meteor. Soc., 549–572, [https://doi.org/10.1007/978-1-935704-20-1\\_23](https://doi.org/10.1007/978-1-935704-20-1_23).

Schemm, S., 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>.

Schultz, D. M., 2005: A review of cold fronts with prefrontal troughs and wind shifts. *Mon. Wea. Rev.*, **133**, 2449–2472, <https://doi.org/10.1175/MWR2987.1>.

—, and C. F. Mass, 1993: The occlusion process in a midlatitude cyclone over land. *Mon. Wea. Rev.*, **121**, 918–940, [https://doi.org/10.1175/1520-0493\(1993\)121<0918:TOPIAM>2.0.CO;2](https://doi.org/10.1175/1520-0493(1993)121<0918:TOPIAM>2.0.CO;2).

—, and G. Vaughan, 2011: Occluded fronts and the occlusion process: A fresh look at conventional wisdom. *Bull. Amer. Meteor. Soc.*, **92**, 443–466, <https://doi.org/10.1175/2010BAMS3057.1>.

—, B. Antonescu, and A. Chiariello, 2014: Searching for the elusive cold-type occluded front. *Mon. Wea. Rev.*, **142**, 2565–2570, <https://doi.org/10.1175/MWR-D-14-00003.1>.

Shafer, J. C., W. J. Steenburgh, J. A. W. Cox, and J. P. Monteverdi, 2006: Terrain influences on synoptic storm structure and mesoscale precipitation distribution during IPEX IOP3. *Mon. Wea. Rev.*, **134**, 478–497, <https://doi.org/10.1175/MWR3051.1>.

Shapiro, M. A., and D. Keyser, 1990: *Extratropical Cyclones: The Erik Palmén Memorial Volume*, Amer. Meteor. Soc., 167–191.

Shen, R., L. Gao, and Y.-A. Ma, 2022: On optimal early stopping: Over-informative versus under-informative parametrization. arXiv, 2202.09885v2, <https://doi.org/10.48550/arXiv.2202.09885>.

Sills, D. M. L., J. W. Wilson, P. I. Joe, D. W. Burgess, R. M. Webb, and N. I. Fox, 2004: The 3 November tornadic event

during Sydney 2000: Storm evolution and the role of low-level boundaries. *Wea. Forecasting*, **19**, 22–42, [https://doi.org/10.1175/1520-0434\(2004\)019<0022:TNTEDS>2.0.CO;2](https://doi.org/10.1175/1520-0434(2004)019<0022:TNTEDS>2.0.CO;2).

Sima, C., and E. R. Dougherty, 2008: The peaking phenomenon in the presence of feature-selection. *Pattern Recognit. Lett.*, **29**, 1667–1674, <https://doi.org/10.1016/j.patrec.2008.04.010>.

Simmonds, I., K. Keay, and J. A. T. 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>.

Simpson, J. E., 1972: Effects of the lower boundary on the head of a gravity current. *J. Fluid Mech.*, **53**, 759–768, <https://doi.org/10.1017/S0022112072000461>.

Sønderby, C. K., and Coauthors, 2020: MetNet: A neural weather model for precipitation forecasting. arXiv, 2003.12140v2, <https://doi.org/10.48550/arXiv.2003.12140>.

Steenburgh, W. J., C. F. Mass, and S. A. Ferguson, 1997: The influence of terrain-induced circulations on wintertime temperature and snow level in the Washington Cascades. *Wea. Forecasting*, **12**, 208–227, [https://doi.org/10.1175/1520-0434\(1997\)012<0208:TIOTIC>2.0.CO;2](https://doi.org/10.1175/1520-0434(1997)012<0208:TIOTIC>2.0.CO;2).

Terven, J., D. M. Cordova-Esparza, A. Ramirez-Pedraza, E. A. Chavez-Urbiola, and J. A. Romero-Gonzalez, 2023: Loss functions and metrics in deep learning. arXiv, 2307.02694v4, <https://doi.org/10.48550/arXiv.2307.02694>.

Thomas, C. M., and D. M. Schultz, 2019: 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>.

Uccellini, L. W., S. F. Corfidi, N. W. Junker, P. J. Kocin, and D. A. Olson, 1992: Report on the surface analysis workshop held at the National Meteorological Center 25–28 March 1991. *Bull. Amer. Meteor. Soc.*, **73**, 459–472.

Wakimoto, R. M., and B. L. Bosart, 2001: Airborne radar observations of a warm front during FASTEX. *Mon. Wea. Rev.*, **129**, 254–274, [https://doi.org/10.1175/1520-0493\(2001\)129<0254:AROOAW>2.0.CO;2](https://doi.org/10.1175/1520-0493(2001)129<0254:AROOAW>2.0.CO;2).

Yang, S., W. Xiao, M. Zhang, S. Guo, J. Zhao, and F. Shen, 2023: Image data augmentation for deep learning: A survey. arXiv, 2204.08610v2, <https://doi.org/10.48550/arXiv.2204.08610>.

Young, G. S., and R. H. Johnson, 1984: Meso- and microscale features of a Colorado cold front. *J. Climate Appl. Meteor.*, **23**, 1315–1325, [https://doi.org/10.1175/1520-0450\(1984\)023<1315:MAMFOA>2.0.CO;2](https://doi.org/10.1175/1520-0450(1984)023<1315:MAMFOA>2.0.CO;2).

Zhang, Y., M. Long, K. Chen, L. Xing, R. Jin, M. I. Jordan, and J. Wang, 2023: Skilful nowcasting of extreme precipitation with NowcastNet. *Nature*, **619**, 526–532, <https://doi.org/10.1038/s41586-023-06184-4>.

Ziegler, C. L., and C. E. Hane, 1993: An observational study of the dryline. *Mon. Wea. Rev.*, **121**, 1134–1151, [https://doi.org/10.1175/1520-0493\(1993\)121<1134:AOSOTD>2.0.CO;2](https://doi.org/10.1175/1520-0493(1993)121<1134:AOSOTD>2.0.CO;2).