

# Machine Learning Estimation of Maximum Vertical Velocity from Radar

RANDY J. CHASE,<sup>a,b</sup> AMY MCGOVERN,<sup>c,d,e</sup> CAMERON R. HOMEYER,<sup>d</sup> PETER J. MARINESCU,<sup>b</sup>  
AND COREY K. POTVIN<sup>d,e,f</sup>

<sup>a</sup> *Cooperative Institute for Research in the Atmosphere, Colorado State University, Fort Collins, Colorado*

<sup>b</sup> *Department of Atmospheric Science, Colorado State University, Fort Collins, Colorado*

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

<sup>d</sup> *School of Meteorology, University of Oklahoma, Norman, Oklahoma*

<sup>e</sup> *NSF AI Institute for Research on Trustworthy AI in Weather, Climate, and Coastal Oceanography, University of Oklahoma, Norman, Oklahoma*

<sup>f</sup> *National Severe Storms Laboratory, Norman, Oklahoma*

(Manuscript received 13 October 2023, in final form 31 January 2024, accepted 8 February 2024)

**ABSTRACT:** The quantification of storm updrafts remains unavailable for operational forecasting despite their inherent importance to convection and its associated severe weather hazards. Updraft proxies, like overshooting top area from satellite images, have been linked to severe weather hazards but only relate to a limited portion of the total storm updraft. This study investigates if a machine learning model, namely, U-Nets, can skillfully retrieve maximum vertical velocity and its areal extent from three-dimensional gridded radar reflectivity alone. The machine learning model is trained using simulated radar reflectivity and vertical velocity from the National Severe Storm Laboratory's convection permitting Warn-on-Forecast System (WoFS). A parametric regression technique using the sinh–arcsinh–normal distribution is adapted to run with U-Nets, allowing for both deterministic and probabilistic predictions of maximum vertical velocity. The best models after hyperparameter search provided less than 50% root mean squared error, a coefficient of determination greater than 0.65, and an intersection over union (IoU) of more than 0.45 on the independent test set composed of WoFS data. Beyond the WoFS analysis, a case study was conducted using real radar data and corresponding dual-Doppler analyses of vertical velocity within a supercell. The U-Net consistently underestimates the dual-Doppler updraft speed estimates by 50%. Meanwhile, the area of the 5 and 10 m s<sup>-1</sup> updraft cores shows an IoU of 0.25. While the above statistics are not exceptional, the machine learning model enables quick distillation of 3D radar data that is related to the maximum vertical velocity, which could be useful in assessing a storm's severe potential.

**SIGNIFICANCE STATEMENT:** All convective storm hazards (tornadoes, hail, heavy rain, straight line winds) can be related to a storm's updraft. Yet, there is no direct measurement of updraft speed or area available for forecasters to make their warning decisions from. This paper addresses the lack of observational data by providing a machine learning solution that skillfully estimates the maximum updraft speed within storms from only the radar reflectivity 3D structure. After further vetting the machine learning solutions on additional real-world examples, the estimated storm updrafts will hopefully provide forecasters with an added tool to help diagnose a storm's hazard potential more accurately.

**KEYWORDS:** Severe storms; Radars/Radar observations; Operational forecasting; Deep learning; Regression

## 1. Introduction

Weather hazards in the United States cost billions of dollars annually (NCEI 2023). A majority of the billion-dollar disaster events involve hazards directly created by convective storms (e.g., hail, tornadoes, floods, straight-line winds). Convective weather hazards are ultimately connected to the fast current of upward moving air, also known as an updraft. Despite the updraft's importance in convective storms, an updraft's intensity or area is not something that is reliably quantified in real time to be used as a forecasting tool for assessing storm severity.

One method to measure a storm's updraft speed is to use multiple Doppler velocity radar measurements of the same storm (see chapter 12.4.3 in Rauber and Nesbitt 2018). Using the continuity equation and the horizontal radial velocity

from the multiple radars, a three-dimensional (3D) wind vector can be derived from which the vertical component is the vertical velocity. While multi-Doppler wind retrievals have been used for decades (e.g., Doviak et al. 1976), the operational ground-based radar system in the United States (i.e., NEXRAD; Crum and Alberty 1993) is not configured in a way that is conducive for high-quality multi-Doppler measurements of storm updrafts (e.g., large baselines, missing low-level coverage). Furthermore, the sensitivity of multi-Doppler analyses to artifacts usually requires some manual quality control by a domain expert, making a timely real-time product prohibitive [see section 3 in Alford et al. (2019) for an example of multi-Doppler quality-control efforts]. Given the lack of direct measurements of storm vertical velocity, several proxy methods have been used in a research capacity to diagnose relationships between other storm updraft characteristics and their severe hazards: overshooting tops (e.g., Marion et al. 2019), lightning (e.g., Carey et al. 2019), differential reflectivity ( $Z_{DR}$ ) columns (e.g., French and Kingfield 2021), bounded

Corresponding author: Randy J. Chase, randy.chase@colostate.edu

DOI: 10.1175/AIES-D-23-0095.1 e230095

© 2024 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)).

Brought to you by NOAA Central Library | Unauthenticated | Downloaded 04/05/24 03:36 PM UTC

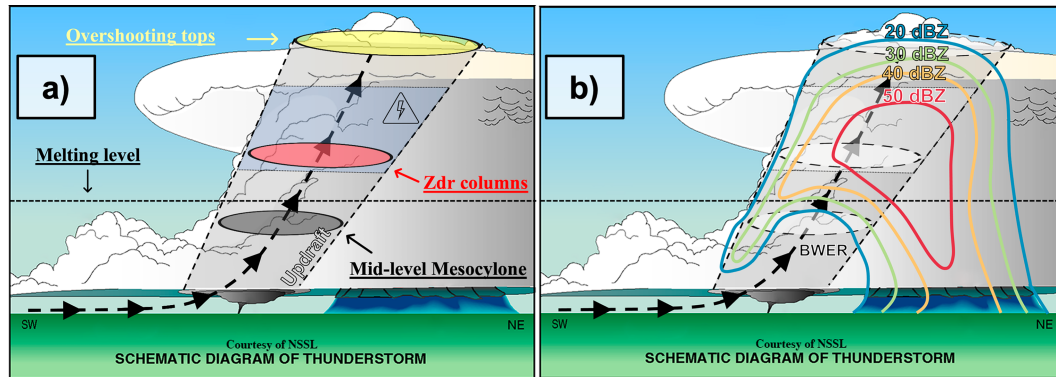


FIG. 1. (a) Schematic summarizing past work on relating storm updraft proxies and their physical location. The black circle represents the approximate location of midlevel mesocyclones (e.g., Sessa and Trapp 2020), the red circle illustrates the typical height of  $Z_{DR}$  columns (e.g., Wilson and Van Den Broeke 2022), the blue shaded region is the approximate region for lightning charging (e.g., Deierling and Petersen 2008), and the yellow circle shows the location of an overshooting top (e.g., Marion et al. 2019). (b) Contours of measured radar reflectivity volumes and how it could outline the full storm updraft [inspired by Chisholm and Renick (1972), reproduced in Trapp (2013)]. The location of a BWER (Browning 1965) is annotated. The background supercell illustration for both (a) and (b) was provided by the NSSL.

weak echo regions (BWER; e.g., Lakshmanan 2000), and mesocyclone width (e.g., Sessa and Trapp 2020).

Overshooting tops are convective updrafts that exceed their level of neutral thermal buoyancy (i.e., equilibrium level), creating an area of colder cloud-top brightness temperatures that can be detected from imaging satellites (e.g., GOES; Bedka et al. 2010). Storms that have overshooting tops have been linked to severe weather on the ground (e.g., Reynolds 1980; Negri and Adler 1981). More quantitatively, the width of the overshooting top [i.e., overshooting top area (OTA)] has been associated with tornado properties like intensity (Trapp et al. 2017) and damage rating (i.e., Marion et al. 2019). Both Trapp et al. (2017) and Marion et al. (2019) found evidence that a wider OTA, and thus a wider overall updraft, is associated with more damaging tornadoes. The key hypothesis from Trapp et al. (2017) and Marion et al. (2019), is that a wider updraft can support a wider tornadic circulation and thus a more damaging tornado. Note that the OTA method can only obtain direct information of the storm's top (Fig. 1a, yellow oval). Thus, the OTA is less informative of the updraft as a whole, including the low-level updraft (i.e., 1 km above ground level), which has been postulated as more important for tornado formation (Peters et al. 2023).

Another method to assess convective updraft characteristics includes lightning, which forms from the release of the buildup of electrical charge originating from rimed ice and pristine ice collisions within a storm's updraft. Deierling and Petersen (2008) and Carey et al. (2019) show that a storm's updraft volume has a clear relationship with total lightning activity. Furthermore, other studies have shown that a lightning jump, a rapid increase in lightning activity, can occur prior to severe weather (e.g., Williams et al. 1999) and can be linked to the overall updraft kinematics (Schultz et al. 2015). While there is evidence of a strong link between updraft volume and total lightning activity, storm microphysics is highly variable

and storm-to-storm differences make building a general quantitative relationship challenging. Furthermore, the best correlation between the updraft volume and lightning activity only pertains to a subset of the total updraft area (Fig. 1a, blue shading).

Beyond overshooting tops and lightning, an additional proxy method for diagnosing information of a storm's updraft is  $Z_{DR}$  columns (e.g., Kumjian and Ryzhkov 2008). In some convective updrafts, storms loft raindrops above the melting level where they retain their semioblate shape that corresponds to a positive  $Z_{DR}$  value (i.e., 1–4 dB) that is surrounded by ice particles that have near-0-dB  $Z_{DR}$ . Thus, the high  $Z_{DR}$  region outlines a portion of the updraft. French and Kingfield (2021) and Wilson and Van Den Broeke (2022) investigated how the  $Z_{DR}$  column area is related to tornadoes, showing evidence that a larger  $Z_{DR}$  column area is correlated with a more damaging tornado. Furthermore, Kuster et al. (2020) showed that  $Z_{DR}$  columns change more clearly and earlier with updraft strength than reflectivity changes. While  $Z_{DR}$  columns seem like a promising operational storm diagnostic, the region of the updraft that is highlighted by the  $Z_{DR}$  column is ambiguous (e.g., Fig. 1a, red oval) and not always present.

Something else that can be observed by radars that are associated with updrafts is the presence of a BWER also known as a weak echo region (WER) or a radar vault (e.g., Browning 1965; Chisholm 1970). A WER is a relative minimum of radar reflectivity in the vicinity of a storm's updraft. The lack of radar echo stems from the fast-moving air containing only small scatterers (i.e., small droplets). If the WER is bounded on both sides by discernible echo, it is then called a BWER. While these storm features have been known for a long time, their quantitative connection to the overall updraft characteristics have not been fully explored. Methods have been developed to automatically detect BWER (e.g., Lakshmanan 2000; Pal et al. 2006; Mahalik et al. 2019) but have not been

leveraged to understand the overall frequency and depth of BWER on an annual or spatial basis. Furthermore, the WER or BWER only characterize the lower portion of the updraft where scatterers are still small (e.g., Fig. 1b). Thus, they only characterize some unknown proportion of the draft and are not always present.

A different radar-based method of measuring updraft characteristics is to measure the width of the mesocyclone. With a rotating updraft, the circulation can be easily observed on Doppler radar, and then metrics like the mesocyclone width (distance between the maximum inbound and outbound velocity) can be extracted. The extracted properties of the mesocyclone have proven operationally useful for distinguishing a tornadoes potential intensity (Gibbs 2016; Gibbs and Bowers 2019; Sessa and Trapp 2020) providing support for the hypothesis that the updrafts characteristics are related to the tornado characteristics put forward by Trapp et al. (2017). Similar to the previously mentioned methods, it is unclear which part of the updraft is being shown by the mesocyclone (e.g., Fig. 1a, black oval), but the method shows some potential in qualitatively measuring updraft characteristics if the storm has a mesocyclone.

The proxies mentioned here for approximating updraft characteristics, velocity, and width, provide incomplete information about the updraft. Furthermore, the direct measurement of updrafts from multi-Doppler measurements are not possible in real time for nowcasting applications. Thus, to fill the dearth of updraft measurements, this paper investigates if a machine learning method [i.e., U-Net (U-Net)] could retrieve updraft intensity and width from reflectivity data alone. More specifically, a machine learning model is trained to translate maps of radar reflectivity to maps of maximum vertical velocity. The main hypothesis is that the full three-dimensional volume of radar reflectivity data contains structures and patterns in it that can be leveraged by machine learning to infer the updraft characteristics (Fig. 1b).

The rest of this paper is structured as follows: section 2 discusses the datasets used in training and evaluating the machine learning method. Section 3 mentions the specifics of the methods used to engineer a dataset usable for machine learning and the machine learning details. Note that section 3 also contains some improvements to the machine learning method used here. Section 4 discusses the primary results of this paper, and section 5 summarizes and concludes the manuscript.

## 2. Data

A goal of this research is to provide a near-real-time estimate of maximum vertical velocity inferred from observed radar data. Ideally, the machine learning method would be trained using high-quality physics-based retrievals of vertical velocity from radar data. While many multi-Doppler wind retrievals exist (e.g., Stechman et al. 2016; Alford et al. 2019; Stechman et al. 2020), multi-Doppler experiments rarely use the same radar frequencies (e.g., Ka, X, or S band), observing platform (e.g., Doppler on Wheels, NOAA P-3 aircraft), and grid spacing (e.g., 500 m, 1 km), which makes obtaining a

large standardized machine learning dataset challenging. As an alternative, this paper trains a machine learning method using convective allowing numerical weather prediction model simulations where reflectivity and vertical velocity exist on the same grid (i.e., image). More specifically, the data used to train the machine learning is from the National Severe Storms Laboratory's Warn-on-Forecast System (WoFS; Stensrud et al. 2009; Jones et al. 2020). After training on synthetic radar data, the machine learning model is evaluated on a case study from The Colorado State University Convective Cloud Outflows and Updrafts Experiment (C<sup>3</sup>LOUD-Ex; van den Heever et al. 2021; Marinescu et al. 2020) where multi-Doppler measurements are available for comparison.

### a. Training domain: NSSL WoFS

The Warn-on-Forecast System is an ensemble of the Weather and Forecasting Research (WRF) Model (36 data assimilation members, 18 forecast members) with a rapid data assimilation cycle (15 mins) and convective allowing horizontal grid spacing (3 km). The ensemble consists of variations of physics parameterizations, namely, the boundary layer scheme, surface physics and the radiation schemes (see Table 2 in Skinner et al. 2018). The variety in schemes and availability of numerous simulations of various severe weather events make the WoFS dataset attractive for training a machine learning model.

Note that NWP simulations of convective updrafts are imperfect. For example, Varble et al. (2014), Marinescu et al. (2016), and Fan et al. (2017) compared NWP simulations of storms to radar estimates of updrafts and the weather models consistently overestimated the magnitude of updrafts. Furthermore, the choice of microphysics scheme can bias simulated reflectivity structures (e.g., Fan et al. 2017; Morrison et al. 2015) since the radar reflectivity is calculated from the bulk hydrometers prediction and assumes Rayleigh scatterers (see Min et al. 2015 for an example of how reflectivity is calculated in WRF). While these are clear limitations of the training dataset used here, the hypothesis is that since the NWP is physics informed (i.e., contains trusted physical equations), the resulting machine learning model trained from NWP data should also be physics informed. While the machine learning model might not be overly accurate (within  $1 \text{ m s}^{-1}$ ), the expectation for this tool is to be well correlated with the true value and be used more as a storm diagnostic (e.g., will this storm produce severe weather) by forecasters.

### b. Transfer domain: GridRad

The trained machine learning model requires radar reflectivity on a regular three-dimensional grid. For the case study evaluation, we choose to use the GridRad Severe dataset (Murphy et al. 2023). GridRad Severe is a three-dimensional gridded radar product created from NEXRAD. The horizontal grid spacing of GridRad is  $0.02^\circ \times 0.02^\circ$  ( $1.5 \text{ km} \times 1.5 \text{ km}$ ) while the vertical grid spacing is 0.5-km vertical resolution up to 7 km above mean sea level (MSL) and 1 km from 7 to 22 km MSL. The three-dimensional analysis is conducted for every 5 min of data and the analysis is run on approximately 100 severe

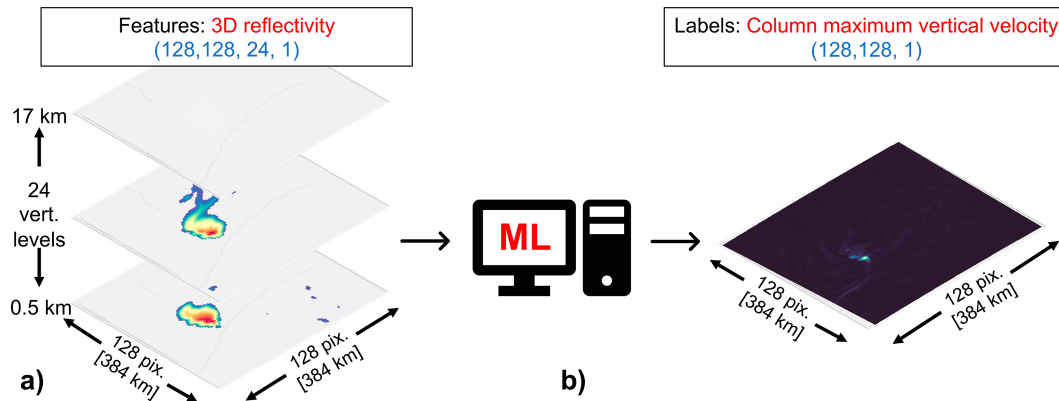


FIG. 2. Training data schematic. (a) The 3D slices of reflectivity used for the input to the machine learning middle (center). (b) Column maximum vertical velocity. Tensor shapes are found in the titles.

weather days per year over the 2010–19 time period. For this paper only 26 May 2017 is used.

### c. Transfer domain: $C^3$ LOUD-Ex data

Out of the  $C^3$ LOUD-Ex field campaign we use 26 May 2017, where dual-Doppler measurements between the Denver WSR-88D KFTG radar and the Colorado State CHILL radar facility (Brunkow et al. 2000) in Greeley, Colorado, are available. More specifically we use the dual-Doppler analysis from Marinescu et al. (2020) where two dual-Doppler techniques were used, namely, Spline Analysis at Mesoscale Utilizing Radar and Aircraft Instrumentation (SAMURAI; Bell et al. 2012) and the Custom Editing and Display of Reduced Information in Cartesian space (CEDRIC; Miller and Fredrick 1998). Both the SAMURAI and CEDRIC dual-Doppler analyses were analyzed on 1-km horizontal grid spacing. For more detailed information on the data used see Marinescu et al. (2020).

## 3. Methods

### a. Dataset engineering

#### 1) WoFS

Several steps are taken to prepare the original WoFS output files for machine learning. The first step is to reduce the number of data that are available (42 TB total of data). For each day that WoFS was run, there are 18 forecast members. For each one of those members, there are several initialization times (e.g., 2200, 2230 UTC), which then have forecast times of 3 or 6 h with 5-min time steps (36 files per initialization time and member). For each day, ensemble member, initialization time and forecast time, one 128 pixel  $\times$  128 pixel  $\times$  24 pixel matrix (384 km  $\times$  384 km  $\times$  17 km) of radar reflectivity data is randomly sliced out of the total domain (Fig. 2a). The dimensions were chosen specifically to allow for at least three maxpooling layers within a U-Net (Ronneberger et al. 2015) architecture (128 is divisible by 2 three times, which is required for U-Net3+). Furthermore, the 384 km  $\times$  384 km is a large enough domain to capture a large portion of storms,

but not too large such that training on the images would be too slow [i.e., larger images require larger random-access memory (RAM) and are slower to train and use]. Other image sizes were not tested here. The corresponding maximum vertical velocity of the column (i.e., the maximum across all heights in WoFS for each grid point) is also extracted on the same 128  $\times$  128 grid to generate the labels (Fig. 2b). Maximum vertical velocity is chosen as a starting point as opposed to translating the three-dimensional radar data to the three-dimensional vertical velocity data, which would be a much more complex task. Only samples where there is at least one pixel of 10 m s<sup>-1</sup> (or greater) are kept so maps without any convection are removed; these maps could otherwise bias the machine learning model toward learning the trivial solution of predicting 0 m s<sup>-1</sup> everywhere.

From the random sliced data, a random subset of examples is chosen for the train, validation, and test set. The exact number of examples are 123 209 (62%), 50 000 (25%), and 25 000 (13%) from years 2019, 2018, and 2020 for the train, validation, and test set, respectively. The year breakdown was selected according to the amount of data that is available: 2019 has the most data, followed by 2018 and 2020 (Fig. 3). Furthermore, saving 2020 for the test dataset enables the most diverse geographic spread of simulations. Note, there is no overlap of data between the train, validation, or test set.

The number of data even after subsetting into the above splits is too large to fit into RAM all at once (e.g., training data are 160 GB, the computers used only have a maximum of 32 GB). Thus, additional efforts are made to make sure the data are only loaded in *as needed* (i.e., on a batch basis). To accomplish this, the reflectivity data is first normalized using a minimum–maximum scaler determined from the training split. Then the entire dataset is converted from float32 precision to float16 precision and saved out to TensorFlow datasets. This enables the use of a dataset larger than available RAM on computing resources available. Furthermore, using TensorFlow datasets anecdotally resulted in higher stability during training and less crashes of the GPU hardware. Hardware and time of computing is mentioned in the following section 4a and Table 1.

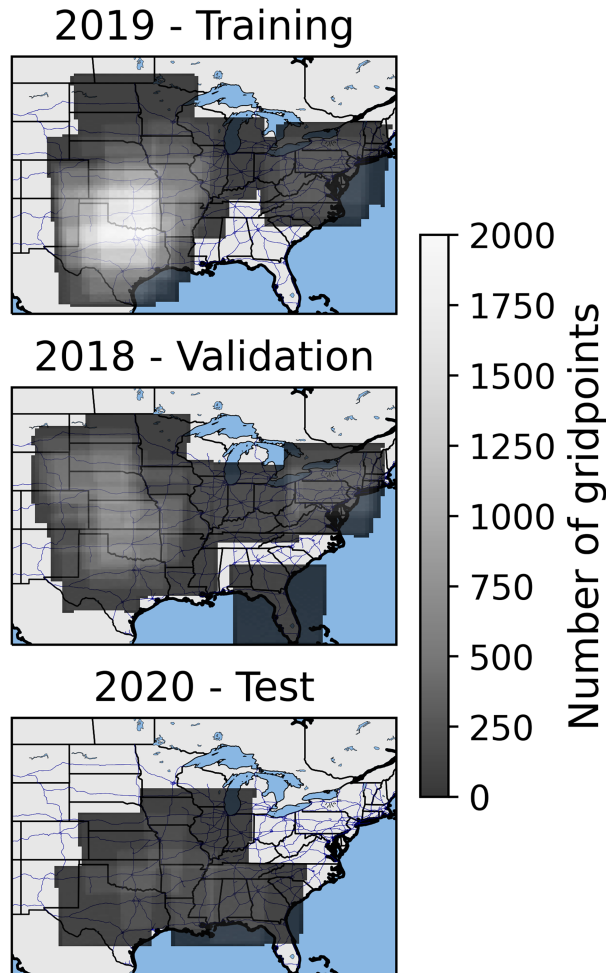


FIG. 3. WoFS domain locations. All domains where WoFS was run in (top) 2019, (middle) 2018, and (bottom) 2020.

## 2) GRIDRAD

Since the machine learning model expects 3-km horizontal spacing images, the GridRad data must be resampled from their original 1.5-km spacing. To do the resampling, a nearest-neighbor and  $k$ -dimensional-tree (kd-tree) approach is employed. The machine learning model also requires the height coordinate to be defined as above ground level. The default height coordinate definition in GridRad is height above mean sea level. While running the GridRad data with mean sea level data likely works well where ground level elevation is close to sea level (i.e., Florida), there were considerable issues

running the model with storms over high terrain (e.g., Wyoming). Thus, to convert from mean sea level to above ground level, the ground elevation at every GridRad point is found using the Shuttle Radar Topography Mission (SRTM) project (Farr et al. 2007). From there, the ground elevation is subtracted from the height profile at each pixel. Then a simple linear interpolation was conducted to get the expected heights the machine learning was trained on (i.e., 0.5–17 km above ground level).

## 3) DUAL-DOPPLER

To make a direct quantitative comparison between the dual-Doppler estimates and the machine learning method all three products were scaled to the same horizontal grid. The process was to first resample the SAMURI dual-Doppler analysis such that it is evaluated on the same exact grid as the CEDRIC dual-Doppler analysis. The resampling is done using a nearest-neighbor approach implemented using kd-trees. From the CEDRIC grid, the data were then smoothed to 3-km horizontal grid spacing, which is the grid spacing of the machine learning product, using the mean. From there, the machine learning method is then resampled from the GridRad grid to the 3-km CEDRIC/SAMURAI grid using a nearest-neighbor approach and kd-trees.

### b. Machine learning method

The main machine learning method used in this paper is U-Net (Ronneberger et al. 2015). U-Nets are well suited for the task of image-to-image translation, which in this paper is the translation of the WoFS simulated reflectivity to the WoFS maximum vertical velocity. Two variants of U-Nets are used in this paper, the standard U-Net (Ronneberger et al. 2015) as well as the upgrade to U-Net, the U-Net3+ (Huang et al. 2020). The main addition to the U-Net3+ method is the addition of full-scale feature connections that add skip connections to all convolutions blocks of the U-Net. Code to build the U-Nets can be found in the data availability section.

Standard loss functions associated with regression tasks include mean squared error or mean absolute error. While these losses can be suitable for some machine learning tasks, they often lead to blurred results when used in image tasks (cf. Fig. 1 in Ravuri et al. 2021). Furthermore, when using mean squared error or mean absolute error there is no source of uncertainty estimate provided. Thus, a parametric regression loss function (e.g., Barnes et al. 2023) is used here.

Parametric regression is where the machine learning model outputs parameters of a distribution [e.g., the mean ( $\mu$ ) and standard deviation ( $\sigma$ ) for the normal distribution] as opposed to deterministic predictions. This enables probabilistic predictions like

TABLE 1. Test dataset statistics. Boldface indicates the best of the four experiments. Asterisks (\*) indicate missing or not calculable. “No. of param.” is the number of trainable parameters for each method.

Models	$R^2$	RMSE ( $\text{m s}^{-1}$ )	$\text{cRMSE}(\text{m s}^{-1}) _{y_{\geq 5}}$	PIT D	IQRr	$\text{IoU} _{y_{\geq 5}}$	No. of param.	CPU (ms)	T4 GPU (ms)
<i>linreg</i>	0.42	1.70	4.80	*	*	0.36	2	*	*
<i>2dmax</i>	0.66	0.75	4.45	<b>0.04</b>	0.62	0.45	203 972	541	<b>58.8</b>
<i>2d24f</i>	0.73	<b>0.67</b>	3.71	<b>0.04</b>	<b>0.52</b>	0.51	207 284	<b>535</b>	147
<i>3d</i>	<b>0.75</b>	0.69	<b>3.67</b>	0.07	<b>0.48</b>	<b>0.51</b>	<b>187 140</b>	9620	702

the median, 75th percentile, and the interquartile range. A recent meteorological example of this technique is [Barnes et al. \(2023\)](#), where parametric regression is used with a multilayer perceptron network for hurricane intensity prediction.

More specifically, [Barnes et al. \(2023\)](#) use a flexible distribution called the sinh–arcsinh–normal distribution (SHASH). The SHASH distribution can be described by four parameters: location ( $\mu$ ), scale ( $\sigma$ ), skewness ( $\gamma$ ), and tailweight ( $\tau$ ). These four parameters are then the four output channels (i.e., maps) of the U-Net that enable the calculation of parameters like the median vertical velocity. The loss function used for the SHASH method is the negative log likelihood:

$$\text{loss} = -\log_e(p), \quad (1)$$

where  $p$  is the value of the SHASH distribution at the target value. We refer readers to [Barnes et al. \(2023\)](#) for more details on the SHASH distribution.

To exemplify this method, consider [Fig. 4](#). Before training, some grid point sent through the neural network might predict the four red parameters, which results in the red curve using the SHASH equation, where the highest likelihood updraft is  $0 \text{ m s}^{-1}$ . The objective of the loss function is to maximize the likelihood of the true value (i.e., the label, which is  $18 \text{ m s}^{-1}$  for this case). After training, the neural network would ideally produce the four parameters in blue that result in the blue curve where the maximum likelihood is now near the observed updraft speed of  $18 \text{ m s}^{-1}$ .

For hyperparameter tuning here, some small adjustments to the SHASH method are made. In the original formulation, the  $\sigma$  and  $\tau$  parameters of the SHASH distribution are positive definite (i.e., cannot be negative). To ensure this property, the output neurons (which are maps here for a U-Net) that represent  $\sigma$  and  $\tau$ ,  $\hat{y}_2$ , and  $\hat{y}_4$ , respectively, are added to the exponent of  $e$ ,

$$\sigma = e^{\hat{y}_2}, \quad \text{and} \quad (2)$$

$$\tau = e^{\hat{y}_4}. \quad (3)$$

While taking the exponent of the output neurons worked for [Barnes et al. \(2023\)](#), often in training the initial parameters of the U-Net, initialized using the glorot uniform initialization, would result in  $\hat{y}_2$  and  $\hat{y}_4$  sufficiently large that when exponentiated become too large to be represented by the machine precision resulting in not a number (NaN) and halted training. An empirical solution to this is to reduce the slope of the exponential relation in Eqs. (2) and (3):

$$\sigma = (e^{\hat{y}_2})^{1/10e}, \quad \text{and} \quad (4)$$

$$\tau = (e^{\hat{y}_4})^{1/10e}. \quad (5)$$

The  $10e$  root power is somewhat arbitrary but is motivated by keeping values roughly in single precision range ( $10^{-6}$ – $10^6$ ).

A second adjustment of the SHASH method is an alteration to the loss. For some instances the likelihood of the label (i.e., truth) was very small (i.e., see the red marker and

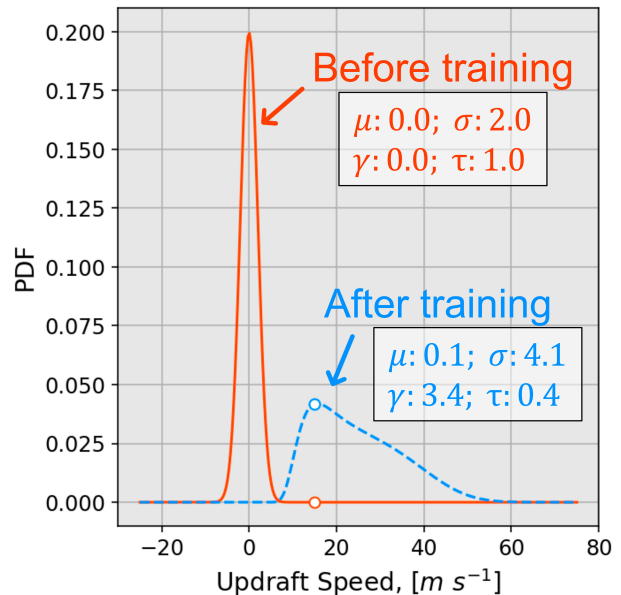


FIG. 4. Training distributions schematic. The red line depicts a grid point at the beginning of training. The circle marker at  $18 \text{ m s}^{-1}$  is the assumed true updraft value for that pixel. The blue line depicts the distribution after training. Example parameter values are noted in the boxes.

line in [Fig. 4](#)), thus the negative log likelihood of that value became large in magnitude. The sum of many large loss values (i.e., the sum of the loss across the images in the batch) can lead to NaN loss values. To combat this issue, a small value ( $\epsilon$ ) is added to the loss equation such that the likelihood of  $p$  never becomes too small:

$$\text{loss} = -\log_e(p + \epsilon), \quad (6)$$

where epsilon is  $10^{-7}$ . The use of  $\epsilon$  is commonly done in machine learning to prevent the division by 0 (i.e., infinity values).

An additional empirical enhancement to the SHASH method comes from the implementation of a weighted loss function. In meteorology, many machine learning tasks of interest are rare (e.g., tornadoes, hail). Even in the context of the number of pixels that have convective updrafts (i.e., greater than  $5 \text{ m s}^{-1}$ ) are far out weighted by pixels near zero. This biases the network toward lower updraft magnitudes. To discourage the model from focusing on the weak updrafts a weighted version of the log likelihood is implemented. Specifically, pixels with an WoFS updraft speed greater than a threshold are weighted differently than those pixels with an updraft speed less than that threshold. Both the weights and the thresholds are static hyperparameters. A weight matrix is created and has the same shape as the loss tensor, where the weight value for any pixel is determined by the WoFS value and the total loss is the element wise product between the weight matrix and the loss tensor. Creating a custom loss function with weighting is inspired by [Ebert-Uphoff et al. \(2021\)](#).

The previous improvements are to the SHASH method directly. It is also important to consider how the data are scaled

before using SHASH. An anecdotal best practice with the SHASH method from this work is to use minimum–maximum scaling. The minimum–maximum scaling forces the input data into a range of 0–1, which results in more reasonable parameter estimates, especially for the first gradient descent updates, compared to other scaling techniques (e.g., standard scalar) and leads to a more stable hyperparameter search.

### c. Hyperparameters

With deep learning the number of hyperparameters can be vast and there are no default parameter sets that are guaranteed to achieve satisfactory results on all machine learning tasks. Thus, a random hyperparameter search is conducted. The following parameters are varied: convolutional kernel size, the number of convolutional filters, the depth of the U-Net, the optimizer, batch normalization, batch size, weight of weighted loss function, the threshold for weighting in the weighted loss function, and the amount of regularization on the convolutional kernels. (For the actual values tested see [Tables A1](#) and [A2](#) in the [appendix](#).) A total of 100 random model configurations are trained per experiment (see next section). All models are allowed to train for up to 200 epochs, but early stopping is set up such that model training ends if overfitting is detected at which point the best weights (smallest loss) are saved. Models rarely trained past 100 epochs, with a majority training for less than 50 epochs.

### d. Model experiment setup

Timeliness is a critically important factor for forecaster tool development ([Harrison 2022](#)). Thus, an experiment is set up to test out three different types of machine learning updraft models with varying complexity and thus inference speeds (experiment names appear in italics throughout). Experiment 1 is to train a 2D convolutional U-Net3+ with the composite reflectivity alone (*2dmax*). The goal behind this model search is to emphasize speed such that the updraft product can be provided to the end user in the fastest time possible. A composite reflectivity input should require the least preprocessing time and the 2D convolution U-Net3+ with one feature should produce the fastest inference time. Experiment 2 is to train a 2D convolutional U-Net3+ on the full 3D reflectivity data [i.e., all 24 levels (*2d24f*)]. This should be the next fastest model for producing predictions. The last model tested is a 3D convolutional U-Net using the full 3D reflectivity data as inputs (*3d*). The full 3D convolutional U-Net should be the slowest. Each one of the three experiments has 100 hyperparameter configurations. The best model from each hyperparameter search is chosen based on the coefficient of determination ( $R^2$ ) between the median of machine learning predicted distribution and the WoFS simulated updraft speed in the validation set.

Beyond the deep learning experiments described above, we also test a simple linear regression model (*linreg*) to confirm the complicated models [i.e., convolutional neural networks (CNNs)] are needed to do the updraft retrieval. It is good practice to compare simple machine learning methods (e.g., [Chase et al. 2022](#)) to more complicated methods (e.g., [Chase et al. 2023](#)) because simpler models are more understandable by end users and therefore tend to be more trustworthy

([Rudin 2019](#)). We train a linear regression on composite reflectivity and only for values greater than 30 dBZ to prevent the linear regression from predicting updrafts for all radar echoes. Because the training and validation datasets are so large, we also train the linear regression off the test set (i.e., 2020). This will give an unfair advantage to the linear regression, but the resulting performance discussed in the next section will show that the linear regression performance is considerably worse than the deep learning experiments even with the advantage of being trained on the test set.

### e. Evaluation metrics

A robust evaluation of machine learning is one that uses multiple metrics. While the best models from the hyperparameter search are determined objectively with  $R^2$ , the additional metrics used for evaluation are the more common statistical definition of  $R^2$  (which is reported in [Table 1](#)), root-mean-squared error (RMSE), conditional root mean squared error (cRMSE), probability integral transform D statistic (PITD), interquartile range hit rate (IQRr), and intersection over union (IoU). RMSE is defined as

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=0}^N (y_i - \hat{y}_i)^2}, \quad (7)$$

where  $y_i$  is the WoFS maximum vertical velocity for some pixel  $i$ ,  $\hat{y}_i$  is the median output from the machine learning method for the same pixel, and  $N$  is the total number of pixels. The cRMSE is defined as

$$\text{cRMSE} = \sqrt{\frac{1}{N} \sum_{i=0}^N (y_i - \hat{y}_i)^2} \Bigg|_{y_i \geq \alpha}, \quad (8)$$

where the equation is the same, but only evaluated for pixels  $i$  that have a  $y_i >$  some threshold  $\alpha$ . IoU is defined as

$$\text{IoU} = \frac{A \cap B}{A \cup B} \Bigg|_{y_i > \alpha}, \quad (9)$$

where  $A$  is the set of pixels where the WoFS maximum vertical velocity is greater than some threshold  $\alpha$ , and  $B$  is the set of pixels where the machine learning median prediction is greater than the same threshold  $\alpha$ . An ideal value of IoU is 1, while values greater than 0.5 are generally considered good. These metrics in general give a numerical value for the pixel-level accuracy (RMSE), the conditional pixel-level accuracy (e.g., for only updraft pixels, how good does the model perform; cRMSE) and how well do the machine learning updrafts overlap with the real updrafts (IoU).

To evaluate the probabilistic information provided by the SHASH method we employ several metrics. The first metric is the IQRr. This is defined as the rate at which the true value (i.e., WoFS updraft maximum) lies within the interquartile range (25th–75th percentile). By definition the IQR should contain 50% of the data, thus an ideal IQRr is 0.5. As a more generalized IQRr, we use additional metrics called the probability integral transform (PIT) histogram and PITD. The PIT

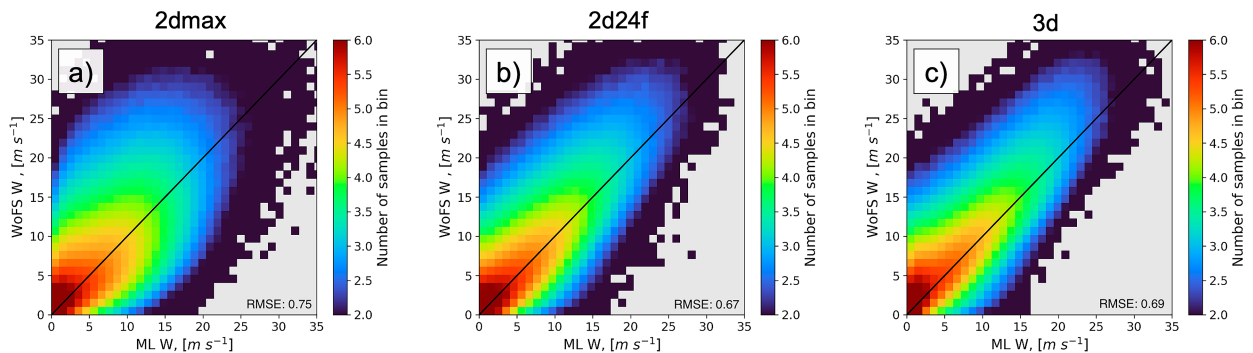


FIG. 5. One-to-one comparison of the median machine learning prediction to the WoFS updraft prediction for the test dataset. The statistics for this comparison can be found in Table 1. (a) Data for the *2dmax* model, (b) data for the *2d24f* model, and (c) data for the *3d* model. The color bar is the log of the counts.

value (used in the histogram) is the quantile of the observed value within the machine learning predicted distribution; in other words, for each quantile range (e.g., 0–0.1), how frequently the observed value falls within this quantile range within the machine learning predicted distributions. Here we binned the PIT values in 0.1 increments (arbitrary choice) and visualized as a normalized histogram. Ideally, the histogram will be flat with all bins having a frequency of the quantile bin width (e.g., 0.1). To quantify how flat the histogram is, the PITD statistic measures the mean deviation from the constant slope, defined as

$$\text{PITD} = \sqrt{\frac{1}{B} \sum_{k=1}^B \left( b_k - \frac{1}{B} \right)^2}, \quad (10)$$

where  $B$  is the number of bins, and  $b_k$  is the relative frequency of the  $k$ th bin. An ideal value of PITD is 0. For more information and examples of the uncertainty metrics please see Haynes et al. (2023) and Barnes et al. (2023).

As an effort of transparency, we also quantify the time it takes to get a machine learning prediction. These timings are solely done on how long it takes to get an updraft prediction from already available and loaded in radar data. We run 30 batches of size 32 through the model and take the mean time. The number of images in the batch (32) is chosen to ensure the machine learning makes an inference on a domain larger than the United States (recall each image is about  $384 \text{ km} \times 384 \text{ km}$ ). The CPU time was clocked on a 2019 Macbook Pro with an i7 processor and 16 GB of RAM while the GPU time was clocked on Google Colab using their freely available T4 GPUs.

## 4. Results and discussion

### a. Bulk model assessment on WoFS data

Training the 300 model configurations (100 per experiment) took more than 500 h of continuous training on a University of Oklahoma–owned computer cluster of NVIDIA A100s. The *2dmax* and *2d24f* models were trained on two A100s while the *3d* was trained on four A100s (40 GB RAM per card). The best hyperparameter set for each experiment is

chosen using the validation set  $R^2$  value. All metrics described in the previous section are calculated using the test set (2020; Fig. 3) and are summarized in Table 1.

Overall, the experiments (*linreg*, *2dmax*, *2d24f*, and *3d*) show expected results, with the deep learning methods performing considerably better than *linreg*. The *3d* and *2d24f* models have very similar deterministic performance and the *2dmax*, which has less input information, has worse results in all metrics except for the PITD statistic. The deterministic statistics for the deep learning models are visualized in Fig. 5, where the pixel wise median from the machine learning predicted distribution is compared directly to the WoFS maximum vertical velocity for the same pixel. If the model prediction was perfect, the output would lie exactly on the diagonal line. The spread about the diagonal is considerably reduced as the model complexity increases (Figs. 5a–c). Meanwhile the uncertainty information visualization (i.e., the PIT Histograms; Fig. 6) does not follow the same trend (Figs. 6a–c). The *2dmax* and *2d24f* models have better uncertainty information (PITD of 0.04) than the *3d* model (PITD of 0.07). Beyond the PITD statistic, the *2dmax* model has a peaked histogram at 0.75 PIT while the other two have more left-skewed distributions, with the probability bins of the largest PIT values exceeding 0.16 for the *3d* model. The interpretation of the peaked shape in Fig. 6a, is that the machine learning model is likely under confident in its predictions, predicting too wide of a distribution. Meanwhile, the skewed-left distributions of Figs. 6b and 6c are interpreted as the true value landing frequencies (more than 10% of the time) within or beyond the upper tail of the machine learning–predicted distribution.

The results of Figs. 5 and 6 suggest that the machine learning predicted distributions are shifted to the left (i.e., closer to 0), which results in the observations being found above the median and at higher quantiles (e.g., 0.75). In general, this indicates an overall low bias in machine learning updraft speed and is supported in the analysis in the following sections. We hypothesize that the low bias is a result of the training data having many pixels where updraft velocities are near  $0 \text{ m s}^{-1}$  with nonzero radar reflectivities. We tried to prevent the underestimation with careful curation of the training dataset [section 3a(1)] and a weighted loss function (section 3b) but



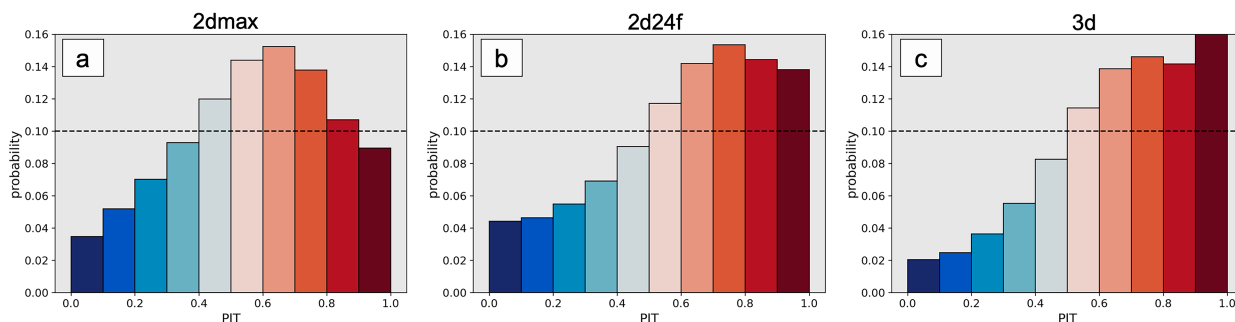


FIG. 6. PIT histograms for the (a)  $2dmax$ , (b)  $2d24f$ , and (d)  $3d$  model evaluated on the test set. The color of the histogram corresponds to the PIT value. The horizontal dashed line is the ideal location of each bar.

we were unsuccessful in removing the low bias completely. In the following sections we show that while there is a low bias, the machine learning updraft model has other beneficial qualities, like providing an estimate on updraft width.

The amount of time for inference on the CPU is about an order of magnitude longer between the 2D convolutions (i.e.,  $2dmax$  and  $2d24f$ ) models and the 3D convolutions (i.e.,  $3d$ ; Table 1). If a 9620-ms runtime is too large for operational uses, there is the potential to use GPUs in the cloud, which would make all methods run in a similar amount of time. If GPUs are not an available resource, then the user could choose to use the  $2d24f$  method, which has similar skill scores to that of the 3D method but at a faster run time.

It is clear that the  $2d24f$  and  $3d$  models are better than the  $2dmax$  model. Distinguishing if the  $2d24f$  or  $3d$  model is

better, is more contested (i.e., show similar statistics in Table 1). For the sake of brevity, the remaining figures in the paper will show the  $3d$  model because of its superior results on the validation set (not shown).

### b. A WoFS example: 30 April 2019

To fully exemplify the machine learning method being used here a case day from the WoFS dataset is used (Fig. 7). Note that this case day (all 18 members and initialization times) is not in the training dataset. On 30 April 2019 a shortwave trough was lifting out from New Mexico and headed toward the Oklahoma–Texas region. Surface-based convective available potential energy (CAPE) values were forecast to be 3000–4000  $J\ kg^{-1}$  over the Texas Panhandle region. Bulk 0–6-km shear was forecast to be around 50 kt (1 kt  $\approx 0.51\ m\ s^{-1}$ ) and storm relative helicity

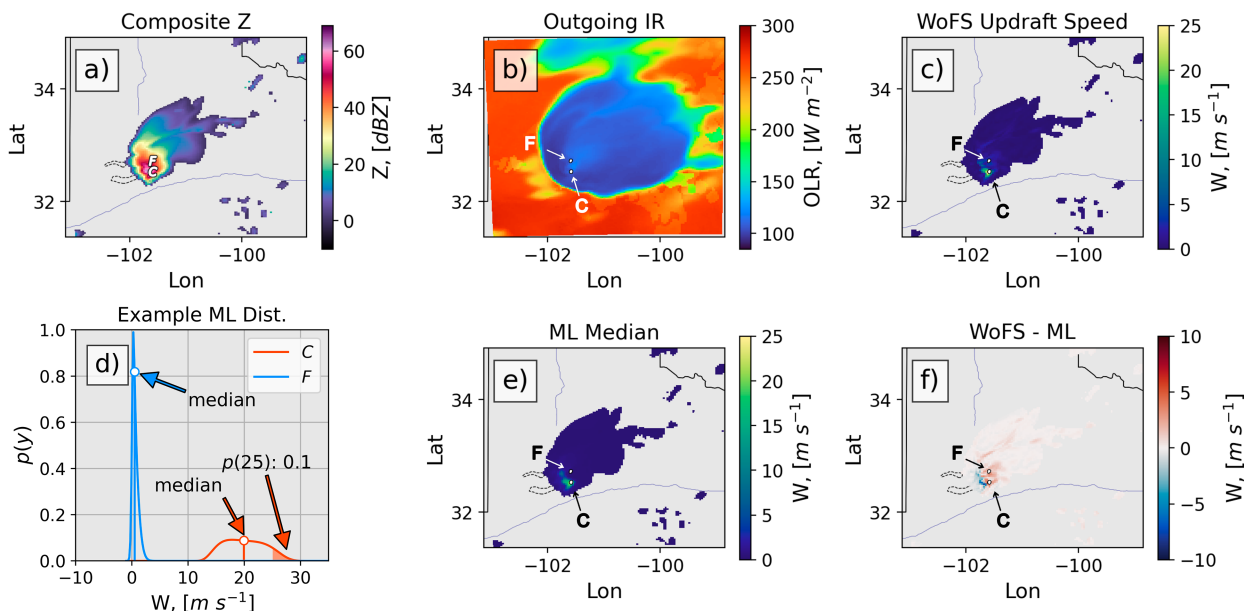


FIG. 7. A WoFS data example from 30 Apr 2019 in north Texas. (a) Composite reflectivity with a marker of “C” for the updraft core and the marker “F” for the forward flank downdraft. The dotted contour is the 1-in. hail location determined by the hail forecast system (HAILCAST; Adams-Selin and Ziegler 2016). (b) Simulated outgoing infrared imagery from the same storm. (c) WoFS simulated maximum vertical velocity masked to where the composite reflectivity is greater than 0. (d) Example machine learning distributions from the C and F markers in (a). The annotation of  $p(25)$  is the integral of the red curve for values greater than 25  $m\ s^{-1}$ . (e) Median maximum vertical velocity from the machine learning model. (f) Difference between maximum vertical velocity from WoFS and the median prediction from the machine learning.

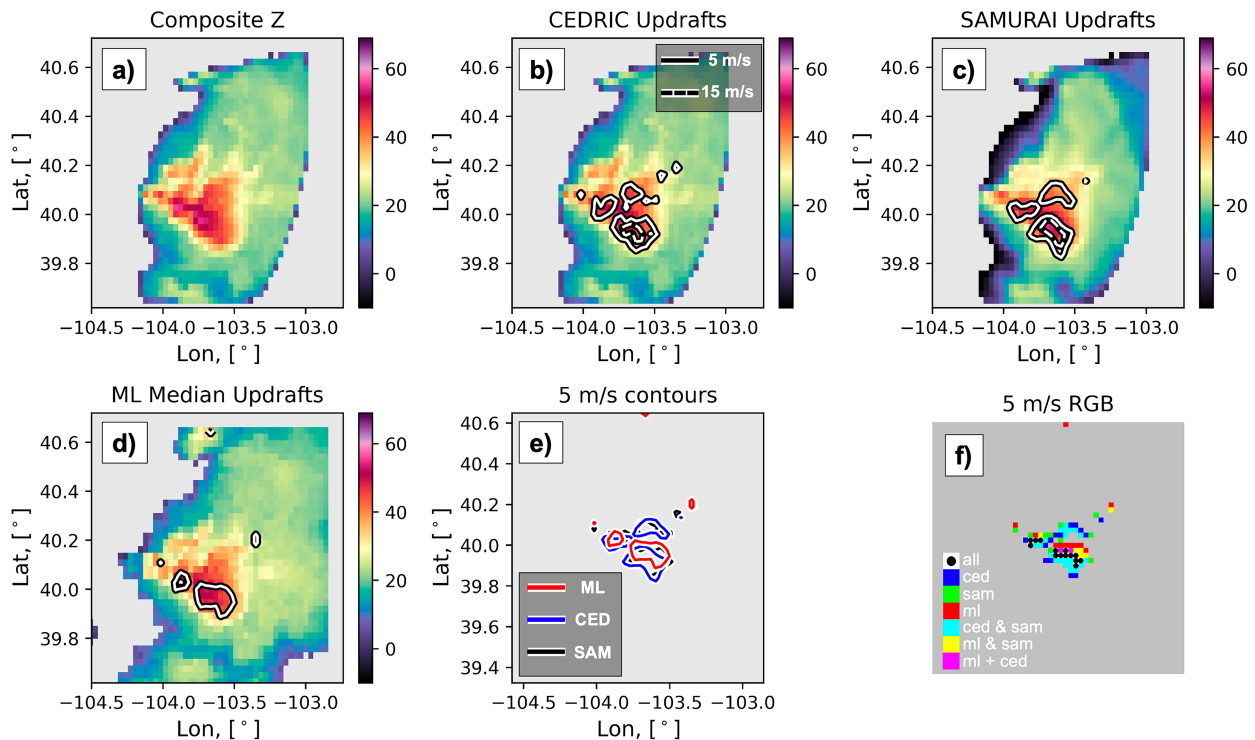


FIG. 8. A supercell case study from eastern Colorado on 26 May 2017. (a) Observed composite reflectivity from the CEDRIC analysis. (b) As in (a), but with column maximum vertical velocity contours of 5 and 15  $\text{m s}^{-1}$ . (c) As in (b), but with SAMURAI. (d) As in (c), but the median machine learning output. (e) All 5  $\text{m s}^{-1}$  contours from CEDRIC (blue), SAMURAI (black), and the machine learning median (red). (f) A red–green–blue (RGB) image where the color value is determined by if the updraft retrieval has more than 5  $\text{m s}^{-1}$  updrafts. The red channel of the RGB image is from the machine learning model, the green channel is from the SAMURAI analysis, and the blue channel is from the CEDRIC analysis. Pixels where all methods overlap are in white with a black circle marker in them.

(SRH) was forecast to be in the 200–350  $\text{m}^2 \text{s}^{-1}$  range, supportive of supercells. By the end of the day multiple hail reports were produced by several supercells that traversed northern Texas between Abilene and Amarillo. The tornadoes that occurred on this day were located in eastern Oklahoma, Missouri, and Arkansas.

WoFS member one, initialized at 2000 UTC, has a single supercell traversing the Texas Panhandle (Fig. 7). Simulated composite reflectivity values exceed 60 dBZ (Fig. 7a), outgoing infrared shows an evident spreading anvil (Fig. 7b) and WoFS simulated maximum vertical velocities exceed 25  $\text{m s}^{-1}$  (Fig. 7c). The median machine learning prediction is plotted in Fig. 7e. Overall, the machine learning model does well, capturing the primary updraft location to within  $\pm 5 \text{ m s}^{-1}$  ( $\sim 20\%$  error) of the WoFS simulated value.

While the median of the machine learning distribution is useful here (Fig. 7e), the full distributions can be interrogated (Fig. 7d). Two locations from within the supercell, the core (C; Fig. 7a) and the forward flank downdraft (F; Fig. 7a), exemplify these distributions. The core updraft location shows a broad distribution with nonzero probability density function (PDF) values from about 12 to 30  $\text{m s}^{-1}$  while the forward flank has a steeply peaked distribution with a peak value just above 0  $\text{m s}^{-1}$ . For this case, the machine learning estimated likelihood of the WoFS simulated updraft maximum (or greater) is 10%. Thus, while the median prediction underestimates the

WoFS updraft, the true value is within the predicted distribution. Another encouraging result is that the likelihood of strong vertical velocities in the forward flank, where downdrafts are expected to dominate, is effectively 0.

### c. Transfer domain case: 26 May 2017

Thus far the comparisons and evaluations of the machine learning model have been restricted to the WoFS data domain. Note that NWP data are not reality and can be considerably different from observations (e.g., Varble et al. 2014; Fan et al. 2017; Marinescu et al. 2016). Furthermore, the motivation of this paper is to train a machine learning model that could be used in real time on observed radar data. Thus, as an out-of-sample test case, the 26 May 2017 dual-Doppler case from Marinescu et al. (2020) is analyzed here.

On this day, terrain-initiated convection formed in an environment with surface-based CAPE above 1500  $\text{J kg}^{-1}$  and SRH greater than 200  $\text{m}^2 \text{s}^{-2}$ , supportive of supercells. A lone supercell formed in the dual-Doppler lobe between KFTG and CHILL, allowing for dual-Doppler syntheses of maximum vertical velocity. The closest analysis time (2144 UTC) with the storm closest to the radars for which the dual-Doppler synthesis errors should be minimized is shown in Fig. 8. Both dual-Doppler analyses (i.e., CEDRIC and SAMURAI) show three broad areas of 5  $\text{m s}^{-1}$  updrafts around the core of reflectivity, with

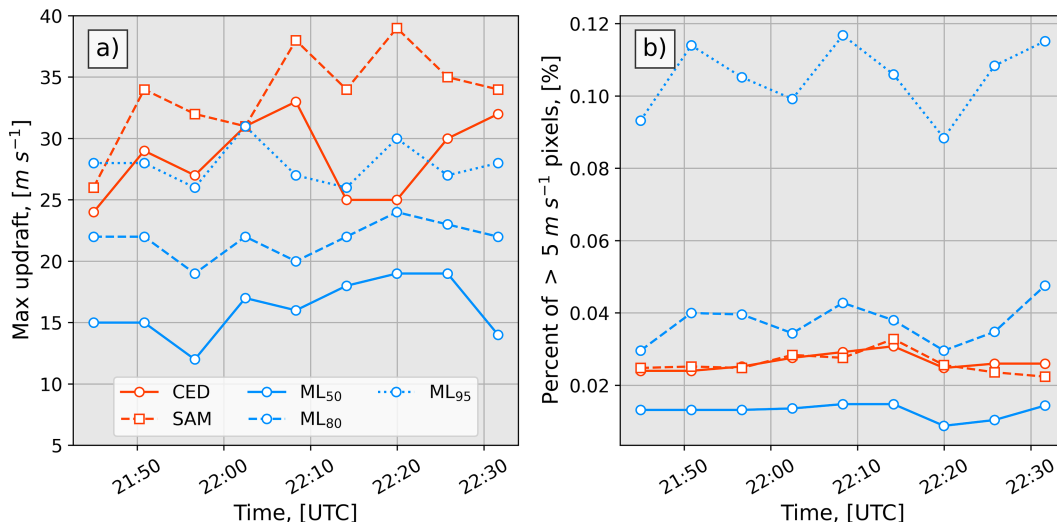


FIG. 9. Updraft statistics for the 26 May 2017 case. (a) Maximum updraft speed for the main updraft through time. Red lines are the dual-Doppler techniques (dashed SAMURAI, solid CEDRIC), while the blue lines are different quantiles of the machine learning distribution (solid median, dashed 80th percentile, and dotted 95th percentile). (b) Percentage of pixels in the entire domain that exceed  $5 \text{ m s}^{-1}$ . Line colors are the same as in (a).

the strongest updrafts found to the south, with values exceeding  $15 \text{ m s}^{-1}$  (maximum of 24 and  $26 \text{ m s}^{-1}$  for CEDRIC and SAMURAI, respectively; Figs. 8b,c). Meanwhile the machine learning median output captures the two southernmost  $5 \text{ m s}^{-1}$  updraft regions. The most intense vertical velocity regions determined by the dual-Doppler are generally not captured in the machine learning median output, and the same southern updraft region only shows the  $5 \text{ m s}^{-1}$  contour (Fig. 8d). The  $5 \text{ m s}^{-1}$  contours from the three analyses partly overlap (14-pixel overlap; Figs. 8e,f). Furthermore, the shapes of the updraft regions, particularly the more elongated shape of the southernmost updraft region is shown in both the dual-Doppler and machine learning data.

Extracting the maximum updraft value for the southernmost updraft over time shows that the median prediction from the machine learning is about half of what the dual-Doppler syntheses are suggesting (Fig. 9a). Given the flexibility of the distribution predictions from the machine learning more than the median prediction can be considered (e.g., 80th percentile and the 95th percentile). For this supercell case the 95th percentile resembles the CEDRIC synthesis better (solid and dotted lines Fig. 9).

Part of the motivation of this work is to see if the machine learning could estimate updraft widths properly because the updraft width is important for some severe weather hazards [e.g., tornadoes and hail in Marion et al. (2019) and Kumjian et al. (2021), respectively]. We use two evaluations of the updraft widths, a simple quantification of updraft pixels and a more complicated IoU statistic. The simple way to quantify the updraft area is to take the percentage of pixels with vertical velocities greater than  $5 \text{ m s}^{-1}$  (out of a total of 2500 pixels; Fig. 9b). Note that taking the area of the column maximum vertical velocity has the added ambiguity for tilted updrafts but is still a first-order estimation of storm updraft area. Similar

to Fig. 9a, the median value of the machine learning distribution is underpredicting the total area of  $5 \text{ m s}^{-1}$ . In contrast to Fig. 9a, where larger percentiles capture the dual-Doppler synthesis, the larger percentiles of the machine learning distribution leads to a systematic overestimation of  $5 \text{ m s}^{-1}$  updraft area.

The more complicated IoU statistic evaluates if the updraft is in the right location. Thus, the IoU is calculated for each time step and several thresholds of updraft speed ( $5$ ,  $10$ , and  $15 \text{ m s}^{-1}$ ) using only the median of the machine learning distribution. The value of IoU as a function of time is shown in Fig. 10. In general, the IoU value between the two dual-Doppler methods is at about 0.5 for all updraft thresholds (Fig. 10a). Meanwhile, the machine learning median predictions shows considerably worse IoU scores compared to CEDRIC (SAMURAI was similar but not shown), where IoU for  $5$  and  $10 \text{ m s}^{-1}$  is on average 0.25 and for  $15 \text{ m s}^{-1}$  IoU can range between 0 and 0.15. The IoU for the machine learning compared to each of the dual-Doppler analyses is about half that of the two dual-Doppler analyses. Since the dual-Doppler techniques have the same input data and follow similar concepts, it is expected that they would have better overlap with each other (Fig. 10a), as compared to the machine learning method (Figs. 10b,c). However, the consistent, albeit smaller, overlap between the machine learning method and dual-Doppler methods, particularly at the weaker vertical velocity thresholds, shows potential skill and application for the machine learning method.

## 5. Conclusions

The severe weather literature notes the importance of a storm's updraft characteristics (i.e., strength and width) in its role in severe weather hazards (e.g., Trapp et al. 2017; Marion et al. 2019; Carey et al. 2019; French and Kingfield 2021;

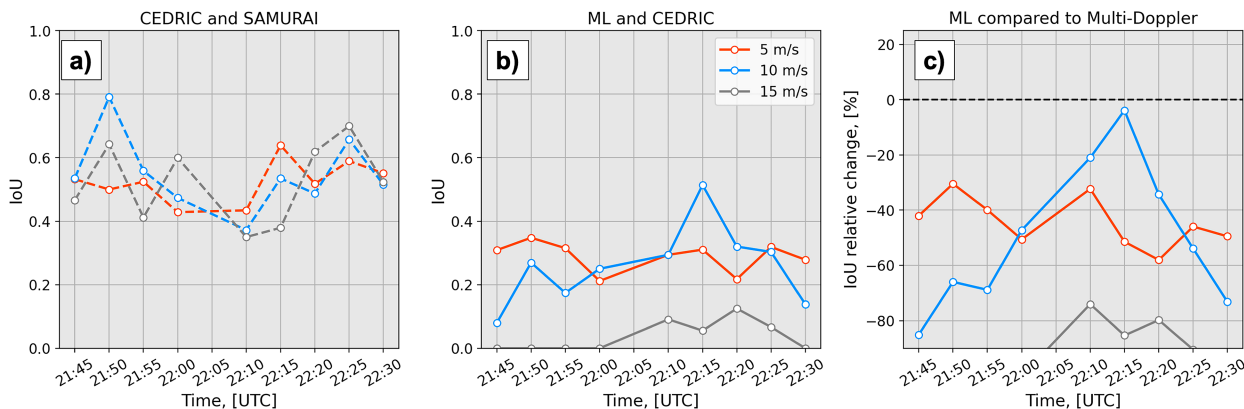


FIG. 10. IoU metric through time. Each plot has a 5, 10, and 15  $\text{m s}^{-1}$  threshold to binarize the data for the IoU calculation (red, blue, and gray, respectively). (a) The two dual-Doppler compared to one another, (b) the median machine learning model output compared to CEDRIC, and (c) the relative difference between (a) and (b).

Kumjian et al. 2021; Wilson and Van Den Broeke 2022). However, only incomplete proxies of storm updraft characteristics exist in real time for forecaster use (see Fig. 1a). Thus, this paper trained a machine learning model, namely, a U-Net (U-Net; Ronneberger et al. 2015; Huang et al. 2020), to translate three-dimensional radar reflectivity data into the column maximum vertical velocity (i.e., the maximum updraft) as an effort to create a near-real-time estimate of storm updraft velocity and area. The U-Net was trained on convective allowing numerical weather prediction data, from the National Severe Storms Laboratory Warn-on-Forecast System (NSSL WoFS; Stensrud et al. 2009; Jones et al. 2020), where the three-dimensional radar data and column maximum vertical velocity are simulated on the same grid (i.e., same image). Beyond training the machine learning model on WoFS data, an in-depth case study using observations from C<sup>3</sup>LOUD-Ex was also conducted. The following are the main contributions and conclusions of the paper:

- 1) We adapted a parametric regression technique from Barnes et al. (2023) to run with a U-Net, and made several training stability enhancements (section 3b).
- 2) We showed that a parametric regression U-Net could skillfully reproduce the WoFS updrafts, having coefficients of determination ( $R^2$ ) of  $>0.65$  (best value of 0.75) and root-mean-squared error (RMSE) on convective updrafts ( $>5 \text{ m s}^{-1}$ ) of less than  $4.5 \text{ m s}^{-1}$  (best value of  $3.67 \text{ m s}^{-1}$ ). Furthermore, the model showed skillful updraft area segmentation, characterized by an intersection over union (IoU) of greater than 0.45 (best value of 0.51; see Table 1 for more results).
- 3) We showed encouraging correspondence between the updraft areas predicted by the machine learning model and those analyzed by two different dual-Doppler techniques in Marinescu et al. (2020). Machine learning updrafts and the dual-Doppler updrafts characterized by IoU values around 0.25, which is half the IoU of the two dual-Doppler techniques compared to one another. However, the machine

learning median updraft magnitude averaged half the dual-Doppler updraft magnitudes.

Overall, the performance of the machine learning model is encouraging given the machine learning model is using only one time step of three-dimensional radar data. A potential improvement would be to include additional input channels to the machine learning. The goal of this paper was to start simply with only one input field (i.e., reflectivity) where the forward model of numerical weather prediction is generally accepted to be representative of real storm structures. Additional input features to the machine learning method could be previous time steps of reflectivity since there is evidence that the temporal evolution of storm reflectivity structures is related to vertical velocity (e.g., Haddad et al. 2022; Prasanth et al. 2023). If forward simulations can be realistically done with dual-polarization parameters from the bulk microphysics of the numerical weather prediction, then differential reflectivity ( $Z_{DR}$ ) and specific differential phase ( $K_{DP}$ ) could also be used as inputs given their vertical structures in storms (e.g., Homeyer et al. 2020, 2023). Alternatively, simulated Doppler radar moments could be included (e.g., radial divergence, azimuthal shear). Another improvement to the technique here would be to retrain the machine learning using higher resolution numerical weather prediction. The use of the 3-km horizontal grid spacing is likely a major limitation of what spatial structures of reflectivity the machine learning is leveraging and what vertical motions are resolved (e.g., Potvin and Flora 2015; Schwartz et al. 2017; Marinescu et al. 2020; Squitieri and Gallus 2020). Thus, models using 1-km grid spacing, which is closer to observed gridded radar resolution and is also being tested for WoFS (Kerr et al. 2023), could assist in the representation of reflectivity distributions as well as further resolving more convective motions. Note that all these additions will add to the computational time required to do machine learning inference, which should not be forgotten since timeliness is a critical aspect of a weather forecasting tool (Harrison 2022).

Beyond direct improvements to the machine learning model, more evaluations should be conducted. One such evaluation would include compiling many multi-Doppler cases that span diverse meteorological conditions (bow echoes, hurricanes, etc.) to contextualize the biases in the existing machine learning model. Another evaluation that could be done is to compare the machine learning output directly to storm reports of severe weather hazards. There might exist some thresholds or key patterns that come before a severe weather hazard that could be used to enhance forecasts. Finally, comparing all the proxy methods and the new machine learning method would highlight the strengths of each method, providing the best practices for each method to overall enhance a forecasters ability to diagnose severe storm potential.

Once the machine learning model is ready for operational use, the transition to operations should be carefully considered. Operational radar data have artifacts like ground clutter and nonmeteorological targets that, when used as input to the nonlinear machine learning model, could lead to spurious updrafts or nonphysical updraft values. Furthermore, observed radar data often miss the lowest levels of storms (Maddox et al. 2002), which could be problematic for a machine learning model that was trained with WoFS where low-level data were available all the time. Despite these caveats, it is encouraging to see the machine learning models perform well on the single test of observational data in this paper, but further tests should be conducted before operational use.

*Acknowledgments.* This material is based upon work supported by the National Science Foundation under Grant ICER-2019758, supporting authors R. J. C. and A. M. We acknowledge the support staff at the OU Supercomputing Center for Education and Research (OSCER) that helped with setting up and maintaining computing facility that enabled this research. Author P. J. M. was provided support by INCUS, a NASA Earth Venture Mission, funded by NASA's Science Mission Directorate and managed through the Earth System Science Pathfinder Program Office under Contract 80LARC22DA011. The C<sup>3</sup>LOUD-Ex field campaign and its associated data used in this study were supported by the Monfort Excellence Fund provided to Susan C. van den Heever as a Monfort Professor at Colorado State University, as well as funding from the National Science Foundation Grant AGS-1409686. C. K. P.'s contribution to this work comprised regular duties at federally funded NOAA/NSSL. We thank William McGovern-Fagg for his help editing the NSSL graphic to make it easier to read. We thank Kayla Hoffman for being the first to work on the project as part of the 2022 research experience for undergraduates at the University of Oklahoma. We thank Daniel Stechman, Stephen Nesbitt, and Matthew Wilson for their valuable input on this project. We also thank the three anonymous reviewers for their kind words and their thorough evaluation of our manuscript.

*Data availability statement.* The data for the test set used here is available for download at <https://doi.org/10.5281/zenodo.10001880>, the dual-Doppler data can be obtained from the

authors of [Marinescu et al. \(2020\)](#). The trained models and scripts can be found on the GitHub repository associated with this manuscript at <https://github.com/ai2es/hradar2updraft>. Unfortunately, the training and validation data are too large to host online (more than 200 GB), thus they are available upon request to the corresponding author. The raw WoFS data are considered experimental and not available for use from the corresponding author. Please reach out to NSSL for more information on the WoFS data.

## APPENDIX

### Hyperparameter Tuning Specifics

All the models shown in the paper are the result of a fairly extensive hyperparameter search. Tables A1 and A2 contain the different hyperparameters that were varied. Note that a total of 100 models were trained for each model type, so it is very possible that not all possible hyperparameter solution sets were run.

TABLE A1. Hyperparameters for the  $2d_{max}$  and  $2d_{24f}$  experiments.

Hyperparameter	Value
No. of convolutional layers	[1]
Convolutional kernel size	[3 × 3, 5 × 5, 7 × 7]
Kernel activation function	[relu]
No. of kernels	[2, 4, 8, 16]
Depth of U-Net	[3, 4, 5]
Optimizer	[adam, rmsprop]
Batchnorm	[on, off]
Batch size	[64, 128, 256, 512]
Loss	[SHASH]
Weight for loss	[1, 2, 3, 4, 5]
Threshold for loss	[0, 1, 5, 10]
L1 Reg	[1 × 10 <sup>-1</sup> , 1 × 10 <sup>-2</sup> , 1 × 10 <sup>-3</sup> , 1 × 10 <sup>-4</sup> , 0.0]
L2 Reg	[1 × 10 <sup>-1</sup> , 1 × 10 <sup>-2</sup> , 1 × 10 <sup>-3</sup> , 1 × 10 <sup>-4</sup> , 0.0]

TABLE A2. Hyperparameters for the  $3d$  experiment.

Hyperparameter	Value
Number of convolutional layers	[1]
Convolutional kernel size	[3 × 3, 5 × 5, 7 × 7]
Kernel activation function	[relu]
Number of kernels	[2, 4, 8, 16]
Depth of U-Net	[1, 2, 3]
Optimizer	[adam, rmsprop]
Batchnorm	[on, off]
Batch size	[64, 128, 256, 512]
Loss	[SHASH]
Weight for loss	[1, 2, 3, 4, 5]
Threshold for loss	[0, 1, 5, 10]
L1 Reg	[1 × 10 <sup>-1</sup> , 1 × 10 <sup>-2</sup> , 1 × 10 <sup>-3</sup> , 1 × 10 <sup>-4</sup> , 0.0]
L2 Reg	[1 × 10 <sup>-1</sup> , 1 × 10 <sup>-2</sup> , 1 × 10 <sup>-3</sup> , 1 × 10 <sup>-4</sup> , 0.0]

## REFERENCES

- Adams-Selin, R. D., and C. L. Ziegler, 2016: Forecasting hail using a one-dimensional hail growth model within WRF. *Mon. Wea. Rev.*, **144**, 4919–4939, <https://doi.org/10.1175/MWR-D-16-0027.1>.
- Alford, A. A., M. I. Biggerstaff, and G. D. Carrie, 2019: Mobile ground-based SMART radar observations and wind retrievals during the landfall of Hurricane Harvey (2017). *Geosci. Data J.*, **6**, 205–213, <https://doi.org/10.1002/gdj3.82>.
- Barnes, E. A., R. J. Barnes, and M. DeMaria, 2023: Sinh-arcsinh-normal distributions to add uncertainty to neural network regression tasks: Applications to tropical cyclone intensity forecasts. *Environ. Data Sci.*, **2**, e15, <https://doi.org/10.1017/eds.2023.7>.
- Bedka, K., J. Brunner, R. Dworak, W. Feltz, J. Otkin, and T. Greenwald, 2010: Objective satellite-based detection of overshooting tops using infrared window channel brightness temperature gradients. *J. Appl. Meteor. Climatol.*, **49**, 181–202, <https://doi.org/10.1175/2009JAMC2286.1>.
- Bell, M. M., M. T. Montgomery, and K. A. Emanuel, 2012: Air-sea enthalpy and momentum exchange at major hurricane wind speeds observed during CBLAST. *J. Atmos. Sci.*, **69**, 3197–3222, <https://doi.org/10.1175/JAS-D-11-0276.1>.
- Browning, K., 1965: Some inferences about the updraft within a severe local storm. *J. Atmos. Sci.*, **22**, 669–677, [https://doi.org/10.1175/1520-0469\(1965\)022<0669:SIATUW>2.0.CO;2](https://doi.org/10.1175/1520-0469(1965)022<0669:SIATUW>2.0.CO;2).
- Brunkow, D., V. N. Bringi, P. C. Kennedy, S. A. Rutledge, V. Chandrasekar, E. A. Mueller, and R. K. Bowie, 2000: A description of the CSU–CHILL national radar facility. *J. Atmos. Oceanic Technol.*, **17**, 1596–1608, [https://doi.org/10.1175/1520-0426\(2000\)017<1596:ADOTCC>2.0.CO;2](https://doi.org/10.1175/1520-0426(2000)017<1596:ADOTCC>2.0.CO;2).
- Carey, L. D., E. V. Schultz, C. J. Schultz, W. Deierling, W. A. Petersen, A. L. Bain, and K. E. Pickering, 2019: An evaluation of relationships between radar-inferred kinematic and microphysical parameters and lightning flash rates in Alabama storms. *Atmosphere*, **10**, 796, <https://doi.org/10.3390/atmos10120796>.
- 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>.
- , —, G. M. Lackmann, and A. McGovern, 2023: A machine learning tutorial for operational meteorology. Part II: Neural networks and deep learning. *Wea. Forecasting*, **38**, 1271–1293, <https://doi.org/10.1175/WAF-D-22-0187.1>.
- Chisholm, A. J., 1970: Alberta hailstorms: A radar study and model. Ph.D. thesis, University of Alberta, 257 pp.
- , and J. H. Renick, 1972: The kinematics of multicell and supercell Alberta hailstorms. *Contribution to Sixth Annual Conf. of the Canadian Meteorological Society*, Toronto, ON, Canada, Canadian Meteorological Society, 24–31.
- Crum, T. D., and R. L. Alberty, 1993: The WSR-88D and the WSR-88D operational support facility. *Bull. Amer. Meteor. Soc.*, **74**, 1669–1688, [https://doi.org/10.1175/1520-0477\(1993\)074<1669:TWATWO>2.0.CO;2](https://doi.org/10.1175/1520-0477(1993)074<1669:TWATWO>2.0.CO;2).
- Deierling, W., and W. A. Petersen, 2008: Total lightning activity as an indicator of updraft characteristics. *J. Geophys. Res.*, **113**, D16210, <https://doi.org/10.1029/2007JD009598>.
- Doviak, R. J., P. S. Ray, R. G. Strauch, and L. J. Miller, 1976: Error estimation in wind fields derived from dual-Doppler radar measurement. *J. Appl. Meteor.*, **15**, 868–878, [https://doi.org/10.1175/1520-0450\(1976\)015<0868:EEIWFDFD>2.0.CO;2](https://doi.org/10.1175/1520-0450(1976)015<0868:EEIWFDFD>2.0.CO;2).
- Ebert-Uphoff, I., R. Lagerquist, K. Hilburn, Y. Lee, K. Haynes, J. Stock, C. Kumler, and J. Q. Stewart, 2021: CIRA guide to custom loss functions for neural networks in environmental sciences—Version 1. arXiv, 2106.09757v1, <https://doi.org/10.48550/ARXIV.2106.09757>.
- Fan, J., and Coauthors, 2017: Cloud-resolving model intercomparison of an MC3E squall line case: Part I—Convective updrafts. *J. Geophys. Res. Atmos.*, **122**, 9351–9378, <https://doi.org/10.1002/2017JD026622>.
- Farr, T. G., and Coauthors, 2007: The Shuttle Radar Topography Mission. *Rev. Geophys.*, **45**, RG2004, <https://doi.org/10.1029/2005RG000183>.
- French, M. M., and D. M. Kingfield, 2021: Tornado formation and intensity prediction using polarimetric radar estimates of updraft area. *Wea. Forecasting*, **36**, 2211–2231, <https://doi.org/10.1175/WAF-D-21-0087.1>.
- Gibbs, J. G., 2016: A skill assessment of techniques for real-time diagnosis and short-term prediction of tornado intensity using the WSR-88D. *J. Oper. Meteor.*, **4**, 170–181, <https://doi.org/10.15191/nwajom.2016.0413>.
- , and B. R. Bowers, 2019: Techniques and thresholds of significance for using WSR-88D velocity data to anticipate significant tornadoes. *J. Oper. Meteor.*, **7**, 117–137, <https://doi.org/10.15191/nwajom.2019.0709>.
- Haddad, Z. S., and Coauthors, 2022: Observation strategy of the incus mission: Retrieving vertical mass flux in convective updrafts from low-Earth-orbit convoys of miniaturized microwave instruments. *IGARSS 2022–2022 IEEE Int. Geoscience and Remote Sensing Symp.*, Kuala Lumpur, Malaysia, Institute of Electrical and Electronics Engineers, 6448–6451, <https://doi.org/10.1109/IGARSS46834.2022.9884264>.
- Harrison, D., 2022: Machine learning co-production in operational meteorology. Ph.D. dissertation, University of Oklahoma, 196 pp.
- Haynes, K., R. Lagerquist, M. McGraw, K. Musgrave, and I. Ebert-Uphoff, 2023: Creating and evaluating uncertainty estimates with neural networks for environmental-science applications. *Artif. Intell. Earth Syst.*, **2**, 220061, <https://doi.org/10.1175/AIES-D-22-0061.1>.
- Homeyer, C. R., T. N. Sandmæl, C. K. Potvin, and A. M. Murphy, 2020: Distinguishing characteristics of tornadic and nontornadic supercell storms from composite mean analyses of radar observations. *Mon. Wea. Rev.*, **148**, 5015–5040, <https://doi.org/10.1175/MWR-D-20-0136.1>.
- , E. M. Murillo, and M. R. Kumjian, 2023: Relationships between 10 years of radar-observed supercell characteristics and hail potential. *Mon. Wea. Rev.*, **151**, 2609–2632, <https://doi.org/10.1175/MWR-D-23-0019.1>.
- 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>.
- Jones, T. A., and Coauthors, 2020: Assimilation of GOES-16 radiances and retrievals into the warn-on-forecast system. *Mon. Wea. Rev.*, **148**, 1829–1859, <https://doi.org/10.1175/MWR-D-19-0379.1>.
- Kerr, C. A., B. C. Matilla, Y. Wang, D. R. Stratman, T. A. Jones, and N. Yussouf, 2023: Results from a pseudo-real-time next-generation 1-km warn-on-forecast system prototype. *Wea. Forecasting*, **38**, 307–319, <https://doi.org/10.1175/WAF-D-22-0080.1>.

- Kumjian, M. R., and A. V. Ryzhkov, 2008: Polarimetric signatures in supercell thunderstorms. *J. Appl. Meteor. Climatol.*, **47**, 1940–1961, <https://doi.org/10.1175/2007JAMC1874.1>.
- , K. Lombardo, and S. Loeffler, 2021: The evolution of hail production in simulated supercell storms. *J. Atmos. Sci.*, **78**, 3417–3440, <https://doi.org/10.1175/JAS-D-21-0034.1>.
- Kuster, C. M., T. J. Schuur, T. T. Lindley, and J. C. Snyder, 2020: Using  $Z_{DR}$  columns in forecaster conceptual models and warning decision-making. *Wea. Forecasting*, **35**, 2507–2522, <https://doi.org/10.1175/WAF-D-20-0083.1>.
- Lakshmanan, V., 2000: Using a genetic algorithm to tune a bounded weak echo region detection algorithm. *J. Appl. Meteor.*, **39**, 222–230, [https://doi.org/10.1175/1520-0450\(2000\)039<0222:UAGATT>2.0.CO;2](https://doi.org/10.1175/1520-0450(2000)039<0222:UAGATT>2.0.CO;2).
- Maddox, R. A., J. Zhang, J. J. Gourley, and K. W. Howard, 2002: Weather radar coverage over the contiguous United States. *Wea. Forecasting*, **17**, 927–934, [https://doi.org/10.1175/1520-0434\(2002\)017<0927:WRCOTC>2.0.CO;2](https://doi.org/10.1175/1520-0434(2002)017<0927:WRCOTC>2.0.CO;2).
- Mahalik, M. C., B. R. Smith, K. L. Elmore, D. M. Kingfield, K. L. Ortega, and T. M. Smith, 2019: Estimates of gradients in radar moments using a linear least squares derivative technique. *Wea. Forecasting*, **34**, 415–434, <https://doi.org/10.1175/WAF-D-18-0095.1>.
- Marinescu, P. J., S. C. van den Heever, S. M. Saleeby, and S. M. Kreidenweis, 2016: The microphysical contributions to and evolution of latent heating profiles in two MC3E MCSs. *J. Geophys. Res. Atmos.*, **121**, 7913–7935, <https://doi.org/10.1002/2016JD024762>.
- , P. C. Kennedy, M. M. Bell, A. J. Drager, L. D. Grant, S. W. Freeman, and S. C. van den Heever, 2020: Updraft vertical velocity observations and uncertainties in high plains supercells using radiosondes and radars. *Mon. Wea. Rev.*, **148**, 4435–4452, <https://doi.org/10.1175/MWR-D-20-0071.1>.
- Marion, G. R., R. J. Trapp, and S. W. Nesbitt, 2019: Using overshooting top area to discriminate potential for large, intense tornadoes. *Geophys. Res. Lett.*, **46**, 12520–12526, <https://doi.org/10.1029/2019GL084099>.
- Miller, L. J., and S. M. Fredrick, 1998: CEDRIC custom editing and display of reduced information in Cartesian space. Tech. Rep., 130 pp., [https://github.com/ai2es/hradar2updraft/blob/main/aux\\_docs/cedric.2009sep\\_doc.pdf](https://github.com/ai2es/hradar2updraft/blob/main/aux_docs/cedric.2009sep_doc.pdf).
- Min, K.-H., S. Choo, D. Lee, and G. Lee, 2015: Evaluation of WRF cloud microphysics schemes using radar observations. *Wea. Forecasting*, **30**, 1571–1589, <https://doi.org/10.1175/WAF-D-14-00095.1>.
- Morrison, H., J. A. Milbrandt, G. H. Bryan, K. Ikeda, S. A. Tessendorf, and G. Thompson, 2015: Parameterization of cloud microphysics based on the prediction of bulk ice particle properties. Part II: Case study comparisons with observations and other schemes. *J. Atmos. Sci.*, **72**, 312–339, <https://doi.org/10.1175/JAS-D-14-0066.1>.
- Murphy, A. M., C. R. Homeyer, and K. Q. Allen, 2023: Development and investigation of GriDrad-severe, a multi-year severe event radar dataset. *Mon. Wea. Rev.*, **151**, 2257–2277, <https://doi.org/10.1175/MWR-D-23-0017.1>.
- NCEI, 2023: U.S. billion-dollar weather and climate disasters. NOAA National Centers for Environmental Information, accessed 15 March 2024, <https://doi.org/10.25921/stkw-7w73>.
- Negri, A. J., and R. F. Adler, 1981: Relation of satellite-based thunderstorm intensity to radar-estimated rainfall. *J. Appl. Meteor.*, **20**, 288–300, [https://doi.org/10.1175/1520-0450\(1981\)020<0288:ROSBTI>2.0.CO;2](https://doi.org/10.1175/1520-0450(1981)020<0288:ROSBTI>2.0.CO;2).
- Pal, N. R., A. K. Mandal, S. Pal, J. Das, and V. Lakshmanan, 2006: Fuzzy rule-based approach for detection of bounded weak-echo regions in radar images. *J. Appl. Meteor. Climatol.*, **45**, 1304–1312, <https://doi.org/10.1175/JAM2408.1>.
- Peters, J. M., B. E. Coffey, M. D. Parker, C. J. Nowotarski, J. P. Mulholland, C. J. Nixon, and J. T. Allen, 2023: Disentangling the influences of storm-relative flow and horizontal streamwise vorticity on low-level mesocyclones in supercells. *J. Atmos. Sci.*, **80**, 129–149, <https://doi.org/10.1175/JAS-D-22-0114.1>.
- Potvin, C. K., and M. L. Flora, 2015: Sensitivity of idealized supercell simulations to horizontal grid spacing: Implications for warn-on-forecast. *Mon. Wea. Rev.*, **143**, 2998–3024, <https://doi.org/10.1175/MWR-D-14-00416.1>.
- Prasanth, S., Z. S. Haddad, R. C. Sawaya, O. O. Sy, M. van den Heever, T. Narayana Rao, and S. Hristova-Veleva, 2023: Quantifying the vertical transport in convective storms using time sequences of radar reflectivity observations. *J. Geophys. Res. Atmos.*, **128**, e2022JD037701, <https://doi.org/10.1029/2022JD037701>.
- Rauber, R. M., and S. W. Nesbitt, 2018: *Radar Meteorology: A First Course*. John Wiley and Sons, 496 pp.
- Ravuri, S., and Coauthors, 2021: Skilful precipitation nowcasting using deep generative models of radar. *Nature*, **597**, 672–677, <https://doi.org/10.1038/s41586-021-03854-z>.
- Reynolds, D. W., 1980: Observations of damaging hailstorms from geosynchronous satellite digital data. *Mon. Wea. Rev.*, **108**, 337–348, [https://doi.org/10.1175/1520-0493\(1980\)108<0337:OODHFG>2.0.CO;2](https://doi.org/10.1175/1520-0493(1980)108<0337:OODHFG>2.0.CO;2).
- Ronneberger, O., P. Fischer, and T. Brox, 2015: U-Net: Convolutional networks for biomedical image segmentation. arXiv, 1505.04597v1, <https://doi.org/10.48550/arXiv.1505.04597>.
- Rudin, C., 2019: Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat. Mach. Intell.*, **1**, 206–215, <https://doi.org/10.1038/s42256-019-0048-x>.
- Schultz, C. J., L. D. Carey, E. V. Schultz, and R. J. Blakeslee, 2015: Insight into the kinematic and microphysical processes that control lightning jumps. *Wea. Forecasting*, **30**, 1591–1621, <https://doi.org/10.1175/WAF-D-14-00147.1>.
- Schwartz, C. S., G. S. Romine, K. R. Fossell, R. A. Sobash, and M. L. Weisman, 2017: Toward 1-km ensemble forecasts over large domains. *Mon. Wea. Rev.*, **145**, 2943–2969, <https://doi.org/10.1175/MWR-D-16-0410.1>.
- Sessa, M. F., and R. J. Trapp, 2020: Observed relationship between tornado intensity and pretornadic mesocyclone characteristics. *Wea. Forecasting*, **35**, 1243–1261, <https://doi.org/10.1175/WAF-D-19-0099.1>.
- Skinner, P. S., and Coauthors, 2018: Object-based verification of a prototype warn-on-forecast system. *Wea. Forecasting*, **33**, 1225–1250, <https://doi.org/10.1175/WAF-D-18-0020.1>.
- Squitiere, B. J., and W. A. Gallus Jr., 2020: On the forecast sensitivity of MCS cold pools and related features to horizontal grid spacing in convection-allowing WRF simulations. *Wea. Forecasting*, **35**, 325–346, <https://doi.org/10.1175/WAF-D-19-0016.1>.
- Stechman, D. M., R. M. Rauber, G. M. McFarquhar, B. F. Jewett, and D. P. Jorgensen, 2016: Interaction of an upper-tropospheric jet with a squall line originating along a cold frontal boundary. *Mon. Wea. Rev.*, **144**, 4197–4219, <https://doi.org/10.1175/MWR-D-16-0044.1>.
- , G. M. McFarquhar, R. M. Rauber, M. M. Bell, B. F. Jewett, and J. Martinez, 2020: Spatiotemporal evolution of the microphysical and thermodynamic characteristics of the 20 June

- 2015 PECAN MCS. *Mon. Wea. Rev.*, **148**, 1363–1388, <https://doi.org/10.1175/MWR-D-19-0293.1>.
- Stensrud, D. J., and Coauthors, 2009: Convective-scale warn-on-forecast system: A vision for 2020. *Bull. Amer. Meteor. Soc.*, **90**, 1487–1500, <https://doi.org/10.1175/2009BAMS2795.1>.
- Trapp, R. J., 2013: *Mesoscale-Convective Processes in the Atmosphere*. Cambridge University Press, 346 pp.
- , G. R. Marion, and S. W. Nesbitt, 2017: The regulation of tornado intensity by updraft width. *J. Atmos. Sci.*, **74**, 4199–4211, <https://doi.org/10.1175/JAS-D-16-0331.1>.
- van den Heever, S. C., and Coauthors, 2021: The Colorado state university convective cloud outflows and updrafts experiment (C<sup>3</sup>LOUD-Ex). *Bull. Amer. Meteor. Soc.*, **102**, E1283–E1305, <https://doi.org/10.1175/BAMS-D-19-0013.1>.
- Varble, A., and Coauthors, 2014: Evaluation of cloud-resolving and limited area model intercomparison simulations using TWP-ICE observations: 1. Deep convective updraft properties. *J. Geophys. Res. Atmos.*, **119**, 13 891–13 918, <https://doi.org/10.1002/2013JD021371>.
- Williams, E., and Coauthors, 1999: The behavior of total lightning activity in severe Florida thunderstorms. *Atmos. Res.*, **51**, 245–265, [https://doi.org/10.1016/S0169-8095\(99\)00011-3](https://doi.org/10.1016/S0169-8095(99)00011-3).
- Wilson, M. B., and M. S. Van Den Broeke, 2022: Using the Supercell Polarimetric Observation Research Kit (SPORK) to examine a large sample of pretornadic and nontornadic supercells. *Electron. J. Severe Storms Meteor.*, **17** (2), <https://ejssm.com/ojs/index.php/site/article/view/85>.