



# Identification of Flux Rope Orientation via Neural Networks

Thomas Narock<sup>1\*</sup>, Ayris Narock<sup>2,3</sup>, Luiz F. G. Dos Santos<sup>4</sup> and Teresa Nieves-Chinchilla<sup>2</sup>

<sup>1</sup>Center for Data, Mathematical, and Computational Sciences, Goucher College, Baltimore, MD, United States, <sup>2</sup>NASA Goddard Space Flight Center, Greenbelt, MD, United States, <sup>3</sup>ADNET Systems Inc., Bethesda, MD, United States, <sup>4</sup>CIRES, University of Colorado, Boulder, CO, United States

Geomagnetic disturbance forecasting is based on the identification of solar wind structures and accurate determination of their magnetic field orientation. For nowcasting activities, this is currently a tedious and manual process. Focusing on the main driver of geomagnetic disturbances, the twisted internal magnetic field of interplanetary coronal mass ejections (ICMEs), we explore a convolutional neural network's (CNN) ability to predict the embedded magnetic flux rope's orientation once it has been identified from *in situ* solar wind observations. Our work uses CNNs trained with magnetic field vectors from analytical flux rope data. The simulated flux ropes span many possible spacecraft trajectories and flux rope orientations. We train CNNs first with full duration flux ropes and then again with partial duration flux ropes. The former provides us with a baseline of how well CNNs can predict flux rope orientation while the latter provides insights into real-time forecasting by exploring how accuracy is affected by percentage of flux rope observed. The process of casting the physics problem as a machine learning problem is discussed as well as the impacts of different factors on prediction accuracy such as flux rope fluctuations and different neural network topologies. Finally, results from evaluating the trained network against observed ICMEs from Wind during 1995–2015 are presented.

**Keywords:** flux rope, neural network, machine learning, space weather, magnetic field

## OPEN ACCESS

### Edited by:

Bala Poduval,  
University of New Hampshire,  
United States

### Reviewed by:

Stefano Markidis,  
KTH Royal Institute of Technology,  
Sweden  
Emilia Kilpua,  
University of Helsinki, Finland

### \*Correspondence:

Thomas Narock  
Thomas.Narock@goucher.edu

### Specialty section:

This article was submitted to  
Space Physics,  
a section of the journal  
Frontiers in Astronomy and Space  
Sciences

**Received:** 17 December 2021

**Accepted:** 01 February 2022

**Published:** 10 March 2022

### Citation:

Narock T, Narock A, Dos Santos LFG  
and Nieves-Chinchilla T (2022)  
Identification of Flux Rope Orientation  
via Neural Networks.  
Front. Astron. Space Sci. 9:838442.  
doi: 10.3389/fspas.2022.838442

## 1 INTRODUCTION

Coronal mass ejections (CMEs) are one of many manifestations of our dynamic Sun. CMEs are responsible for the transport of large quantities of solar mass into the interplanetary medium at very high speeds and in various directions. CMEs are commonly referred to as interplanetary coronal mass ejections (ICMEs) after leaving the solar atmosphere and reaching the interplanetary medium. ICMEs are the main drivers of geomagnetic activity at Earth as well as at other planets and spacecraft throughout the heliosphere (Baker and Lanzerotti, 2008; Kilpua et al., 2017a). *In situ* observations of ICMEs frequently find them to have a combination of an increase in magnetic field strength, low proton plasma temperature,  $\beta_{\text{plasma}}$  below 1, and monotonic rotation of the magnetic field components (Burlaga, 1988). These characteristics are commonly referred to as a Magnetic Cloud (MC) (Burlaga et al., 1981; Klein and Burlaga, 1982). CME eruption theories (Vourlidas, 2014) suggest that a twisting internal magnetic signature—referred to as a flux rope—is always present. While commonly observed, not all ICMEs show the signatures of an internal structure characterized by a flux rope, perhaps resulting from changes during interplanetary evolution (Jian et al., 2006; Manchester et al., 2017). Yet, flux ropes are sufficiently prevalent that they can aid in

space weather forecasting. The observed magnetic field profile depends on a flux rope's orientation and where the spacecraft traverses the structure. The latitudinal and longitudinal deflections of CMEs happen in the lower corona and are not expected to change greatly throughout the interplanetary medium. If flux rope orientation and the spacecraft's crossing trajectory can be determined early enough, this can lead to advanced forecasting as the remaining portion of the flux rope's magnetic field structure can be inferred from physics-based models. The flux rope's internal magnetic field structure is prone to couple with Earth's upper magnetosphere triggering magnetic reconnection processes and allowing the injection of solar magnetic energy into the magnetospheric system. Orientation determines the magnetic field profile observed at Earth and, thus, the geo-effectiveness of the flux rope making early determination of a flux rope's orientation a vital requisite for space weather forecasting. A major challenge to developing such a forecasting system is that information about the internal magnetic structure of ICMEs is often limited to 1D observations of a single spacecraft crossing the structure. This leaves a considerable amount of uncertainty about the three-dimensional structure of the ICME.

Various physics-based flux rope models exist [for example, Lepping et al. (1990) and Nieves-Chinchilla et al. (2019)] that can be used to reconstruct the internal ICME magnetic configuration and provide information on orientation, geometry, and other magnetic parameters such as the central magnetic field. Recent *in situ* observations (Kilpua et al., 2017b; Nieves-Chinchilla et al., 2018; Nieves-Chinchilla et al., 2019; Rodríguez-García et al., 2021), and references therein] are continuing to complement earlier studies (Gosling et al., 1973; Burlaga et al., 1981; Klein and Burlaga, 1982) and enhance our understanding of ICMEs, MCs, and flux ropes. Meanwhile, an increase of space- and ground-based data availability has led to more interest in applications of machine learning within the space weather community [see (Camporeale, 2019), and references therein]. Nguyen et al. (2018) have explored machine learning techniques for automated identification of ICMEs and dos Santos et al. (2020) used a deep neural network to create a binary classifier for flux ropes in the solar wind, determining whether a flux rope was or was not present in a given interval. Recently, Reiss et al. (2021) use machine learning to predict the minimum  $B_z$  value as a magnetic cloud was sweeping past a spacecraft.

We aim to assess a neural network's ability to predict a flux rope's orientation after an ICME is identified. This work is an attempt to understand if a neural network can predict a flux rope's orientation having only seen a portion of the event. If the full magnetic field profile of the flux rope can reliably be reconstructed when the spacecraft is only partially through the flux rope this can provide advanced warning of impending geomagnetic disturbance. Yet, as machine learning is relatively new to space weather, the accuracy of these forecasts, and more generally, which neural network topologies to utilize, are unclear. We begin with a set of exploratory experiments to quantify the capabilities of neural networks in this regard. The results of these experiments then serve as a baseline to begin exploring forecasting.

Here, we extend the binary classifier work of dos Santos et al. (2020) and explore a neural network's ability to predict the orientation, impact parameter, and chirality of an already identified flux rope. We extend the capabilities presented in Reiss et al. (2021) by reconstructing the entire three dimensional magnetic field profile. The neural network is trained using simulated magnetic field measurements over a range of spacecraft trajectories and flux rope orientations. Moreover, we report on the prediction accuracy of the neural network as a function of percentage of flux rope observed. To connect this proof of concept to its potential for real-world use, we also present results from evaluating the neural network on flux ropes observed by the Wind spacecraft. In performing these experiments, we highlight the multiple ways in which this space weather forecasting problem can be cast as a machine learning application and the implications those choices have on prediction accuracy.

In **Section 2** we present our methodology. We describe the flux rope analytical model and the generation of our synthetic data set. **Section 2** also details our neural network designs and training process. **Section 3** presents our results first from the full duration synthetic flux ropes, then from partial duration flux ropes, and ultimately from application to flux ropes observed from the Wind spacecraft. We present a discussion of these results in section 4 along with concluding remarks.

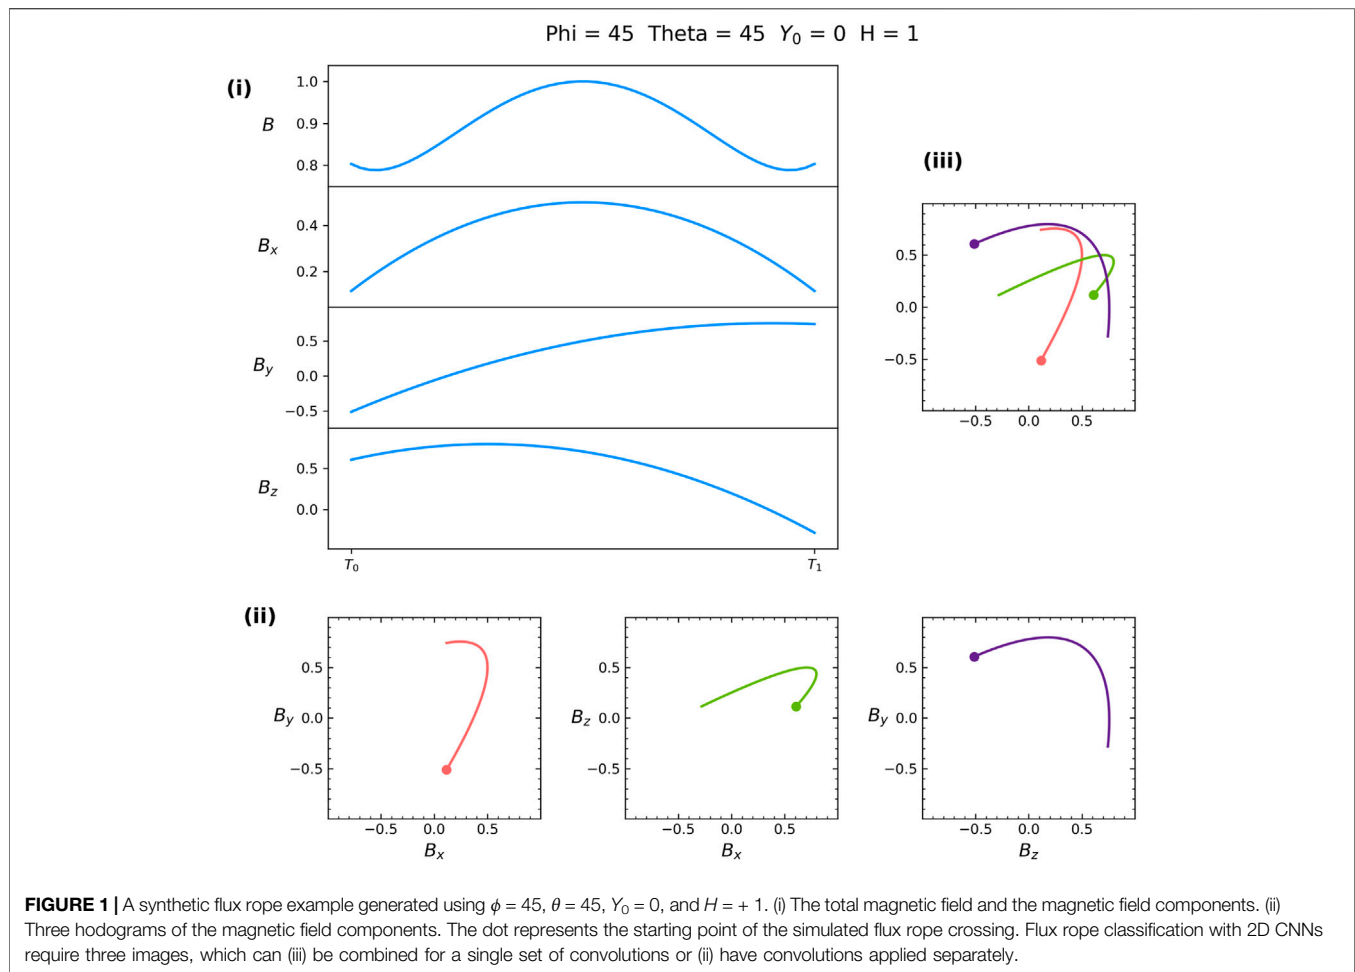
## 2 METHODOLOGY

The task of predicting a flux rope's key defining parameters from magnetic field measurements can be cast as a supervised machine learning problem. This is an approach in which the goal is to learn a function that maps an input to an output based on numerous input-output pairs. There are currently not enough *in situ* observed flux ropes (inputs) with known key parameters (outputs) to train a neural network. Instead, we choose to use a physics-based flux rope model to produce a synthetic training dataset.

### 2.1 Synthetic Data

The circular-cylindrical flux rope model of Nieves-Chinchilla et al. (2016) (N-C model) is used to simulate the magnetic field signature of flux ropes at numerous orientations and spacecraft trajectories. The N-C model takes as input the following parameters:

- $H$ , Chirality of the flux rope; right-handedness is designated with 1, left-handedness with -1
- $Y_0$ , Impact parameter; The perpendicular distance from the center of the flux rope to the crossing of the spacecraft expressed as a percentage of the flux rope's radius
- $\phi$ , Longitude orientation angle of the flux rope
- $\theta$ , Latitude orientation angle of the flux rope
- $R$ , Radius of flux rope
- $V_{sw}$ , Bulk velocity of the solar wind
- $C_{10}$ , A measure of the force free structure. A value of 1 indicates a force free flux rope



The output of the N-C model is the magnetic field profile ( $B_x$ ,  $B_y$ ,  $B_z$  components) that would be observed for spacecraft traversing a flux rope with the given input parameters. An illustration of this is shown in **Figure 1** where panel (i) shows the N-C model output visualized as a time series and panel (ii) depicts the same output as hodograms. All flux ropes were simulated using a solar wind speed ( $V_{sw}$ ) of  $450 \text{ km/s}$ , a radius ( $R$ ) of  $0.07 \text{ AU}$ , and with poloidal normalization. The  $C_{10}$  parameter was held constant at 1, which imposes a force free structure. The model was run for all combinations of longitude ( $\phi$ )  $\in [5^\circ, 355^\circ]$ , latitude ( $\theta$ )  $\in [-85^\circ, 85^\circ]$ , and impact parameter ( $Y_0$ )  $\in [0\%, 95\%]$ . This is done in  $5^\circ$  and  $5\%$  increments and with both chirality options,  $H \in \{-1, 1\}$ . We exclude combinations involving  $\phi = 180^\circ$  as the model is not always defined in this instance. This results in 98,000 combinations.

The fixed bulk velocity of  $450 \text{ km/s}$  and fixed radius of  $0.07 \text{ AU}$  describe a “typical” flux rope observed at Earth based on fittings in Nieves-Chinchilla et al. (2019). Magnetic field profiles of this “typical” flux rope have been shown (dos Santos et al., 2020) to scale with changes in speed and size. In other words, magnetic field profiles are very similar when orientation is held constant and speed and radius are varied. The only variation in the profiles

is duration, which is not a factor for us as all flux ropes are interpolated to 50 points. This relationship allows us to only simulate a subset of all possible speeds and sizes drastically reducing the training data set size and minimizing training time.

The output from each of these 98,000 combinations is then used to generate 10 exemplars of this event in different percentages of completion - from 10 to 100% in steps of 10%. For example, first a 50-point trace through a flux rope defined by the parameter combination is generated (100% completion, **Figure 1(i)**). The first 5 points are interpolated to 50 points to create the 10% completion exemplar. Similarly the first 10 points are used to create the 20% exemplar, the first 15 points for the 30% exemplar, etc. The final dataset contains 980,000 exemplars - a mixture of full duration and partially observed events. These simulated partial flux ropes are useful to understand how much of the flux rope needs to be observed before reliable autonomous predictions can be made. The ability to predict in the absence of the complete flux rope is very desirable in the context of space weather forecasting.

## 2.2 Convolutional Neural Networks

Simply put, a convolution is the application of a filter to an input that results in an activation. Repeatedly applying the same filter to

an input—for example, by sliding a small dimensional filter across an image - results in a map of activations called a feature map. The feature map then indicates the locations and strength of a detected feature in the input. Convolutions are the major building blocks of convolutional neural networks (CNNs) (LeCun and Bengio, 1995), which use a training dataset to learn a set of highly specific filters from the input that lead to the most accurate output predictions. The innovation of the CNN is in not having to handcraft the filters, but rather automatically learning the optimal set of filters during the training process.

The CNN is the basis of the neural network architectures explored in this work. The training phase consists of showing the network the input-output pairs of simulated flux rope magnetic field vectors (input) and the corresponding key parameters used to create this simulated data trace (output). The key parameters represented in this training are  $\phi$ ,  $\theta$ ,  $Y_0$ , and  $H$ . From repeated exposure to input-output pairs the network learns the filters that lead to the most optimal predictions. These neural networks require all inputs to be of the same size, which does not pose a problem when working with synthetic data. *In situ* observations from spacecraft, however, reveal a diverse set of events ranging from a few hours to multiple days. These need to be thoughtfully processed for use as input to the CNN. One could average or interpolate *in situ* events to ensure all input magnetic field time series are of the same length. Alternatively, dos Santos et al. (2020) showed an innovative technique of representing flux ropes as hodograms. Flux ropes of any duration can be cast as a set of three consistently sized images (see **Figure 1(ii)**), which can then serve as input to a CNN. This technique also leverages a wide swath of existing literature in the computer vision field (particularly in the area of handwritten digit classification) that can be helpful in fine-tuning the CNN architecture. Over the next several sections, we present a series of experiments evaluating multiple CNN architectures. Specifically, we compare the predictions from convolutions applied directly to magnetic field time series to predictions made from convolutions applied to hodograms of those magnetic field time series. We do so under two scenarios. First, we develop a baseline for a CNN's capability to predict flux rope orientation by training the architectures with only the 98,000 exemplars of full duration flux ropes. We separately train another copy of each of the three aforementioned architectures with the complete set of 980,000 full and partial duration flux ropes to assess CNN usage in a time-predictive capacity.

### 2.2.1 CNN Architectures

Representing flux ropes as hodograms was inspired by work in handwritten digit classification (dos Santos et al., 2020). Yet, flux ropes provide a more challenging version of this computer vision problem. The input for handwritten digit classification is always a single image; however, flux ropes require a set of three images (hodograms) to capture the entirety of their magnetic field configurations. An initial research question is then how to feed three hodograms as input to a CNN. In the approach chosen for our first architecture, we stack the images (**Figure 1(iii)**) and do a single two-dimensional convolution across the resulting tensor. In our second tested architecture, we apply two-dimensional convolutions to each of the three hodograms separately (**Figure 1(ii)**) and then concatenate the resulting feature maps.

The architecture schematic for the stacked approach is shown in **Figure 2(i)**. An input layer of dimension [100, 100, 3] passes through two rounds of 2D Convolution with a  $3 \times 3$  kernel size. The resulting layer of dimension [100, 100, 64] undergoes a  $2 \times 2$  Max Pooling to transform to dimensions [50, 50, 64]. This layer is then Flattened and Fully Connected to each of four output layers. This 2D CNN with one input ends up with 979,398 trainable parameters.

The architecture for applying two-dimensional convolutions to each of the three hodograms separately and then concatenating the resulting feature maps is shown in **Figure 2(ii)**. Each prong of the initial part of this network involves the same transformations as in the previously described network, with the exception that each of the three input layers is of dimension [100, 100, 1]. Additionally, the Flattened layers at the end of these individual pipelines are then concatenated before being Fully Connected to the four output layers. This architecture has the advantage that salient features in specific hodograms can become more apparent in the feature maps. Yet, this comes at the cost of a more complex neural network. With 2,936,454 trainable parameters, this CNN has significantly more weights that need training.

Finally, we tested an architecture that did not rely on hodogram images. Instead we apply 1D convolutions directly to magnetic field time series. This approach is depicted in **Figure 2(iii)** and results in the smallest CNN with 216,518 trainable parameters. The input layer of size [1, 50, 3] has a 1D Convolution with kernel size 5 applied twice, resulting in a layer of dimension [1, 50, 64]. Max Pooling with a kernel size 2 then creates a layer of dimension [1, 25, 64] before this is flattened to a vector of size 1,600. This layer is then Fully Connected to a layer of size 128 and then to each of the four output layers.

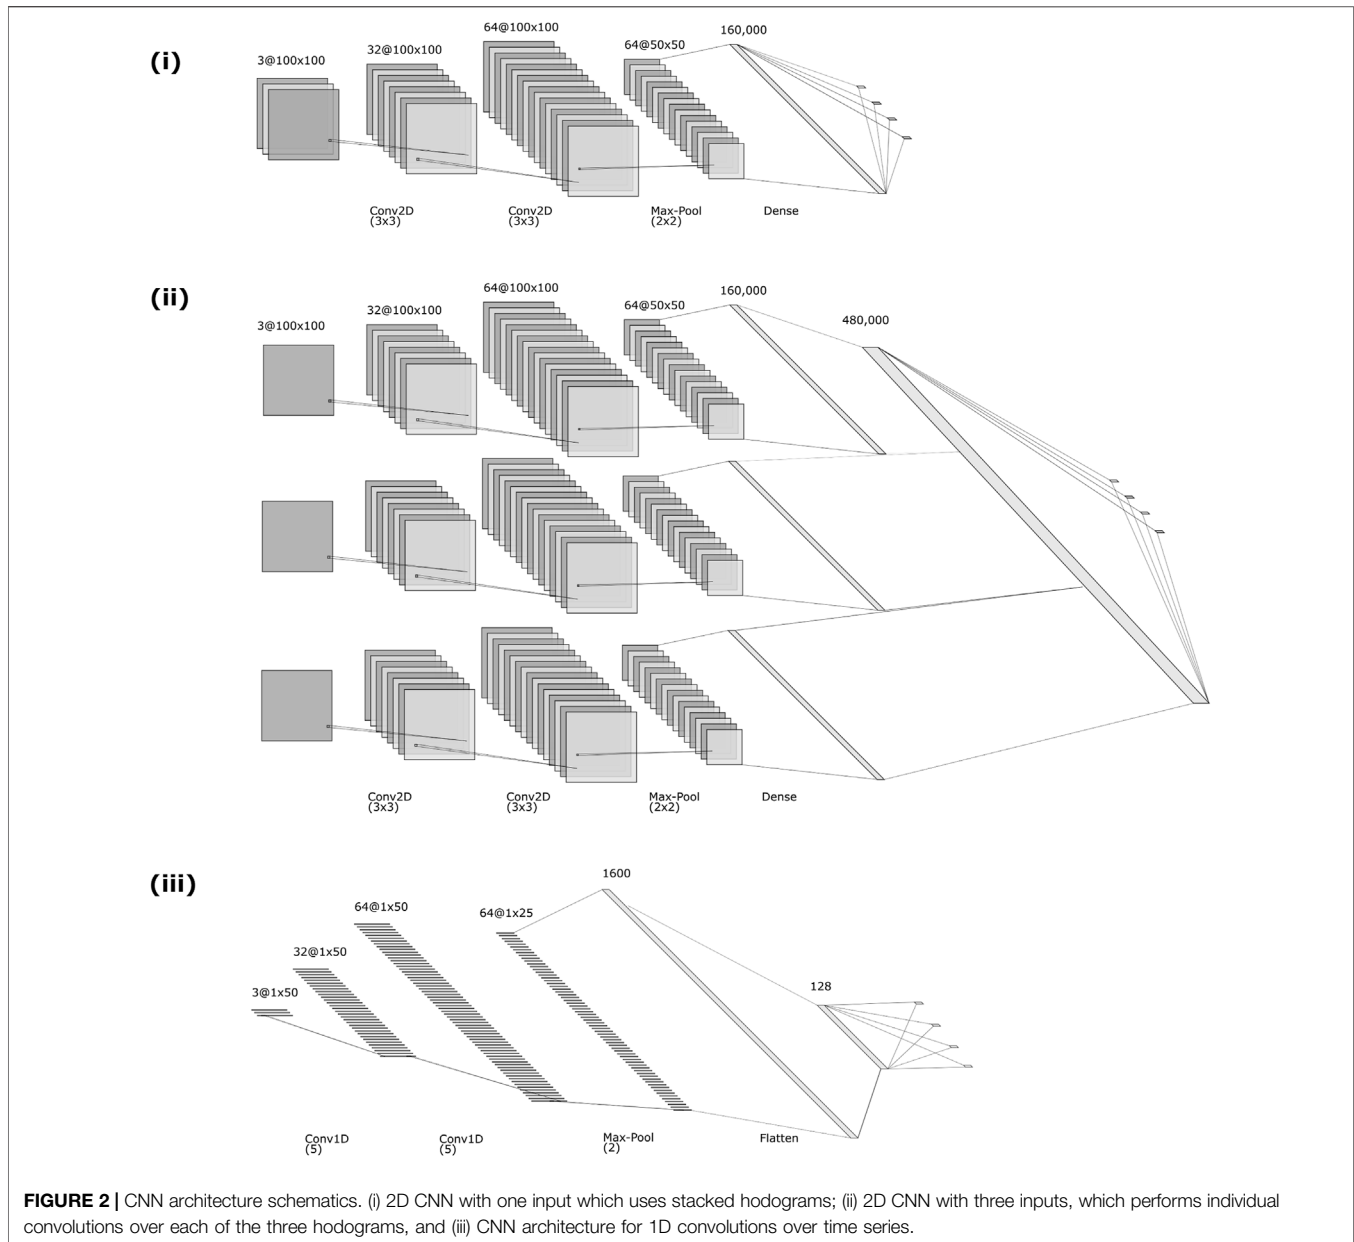
Hyperparameters for all of these architectures were found by doing a simple grid search. Our focus was on comparison of architectures and we acknowledge there may still be room for hyperparameter optimization.

### 2.2.2 CNN Tuning and Training

Neural networks learn by minimizing a loss function, which typically involves some measure of difference between current predictions and expected outputs. Angles can challenge neural network predictions in that loss functions, such as mean squared error (MSE), completely miss the circular nature of angles. For example, if a flux rope's longitudinal value is  $0^\circ$ , then predictions of  $350^\circ$  and  $10^\circ$  are both off by  $10^\circ$ . Yet, MSE will miss this relation and penalizes the  $350^\circ$  prediction more than the  $10^\circ$  prediction. To combat this, we predict  $(\sin(\chi), \cos(\chi))$  with tanh activation to enforce outputs to be in  $[-1, +1]$ . We then post-process the CNN's predictions with arctan to convert to degrees. This approach is applied across all three CNN architectures when predicting  $\phi$  and  $\theta$ .

A challenge also arises in that predicting the real-valued parameters  $\phi$ ,  $\theta$ , and  $Y_0$  is a regression problem while determining the binary parameter,  $H$ , is a classification problem. We address this by training four separate loss functions in each CNN. For  $\phi$  and  $\theta$  we predict the pair  $(\sin(\chi), \cos(\chi))$  and train using the MSE loss function. Impact parameter is also trained using MSE while chirality is defined as a two class classification problem and trained using binary cross entropy.

In our first experiment, the 98,000 full duration synthetic flux ropes were randomly divided into 60% training, 20% validation,



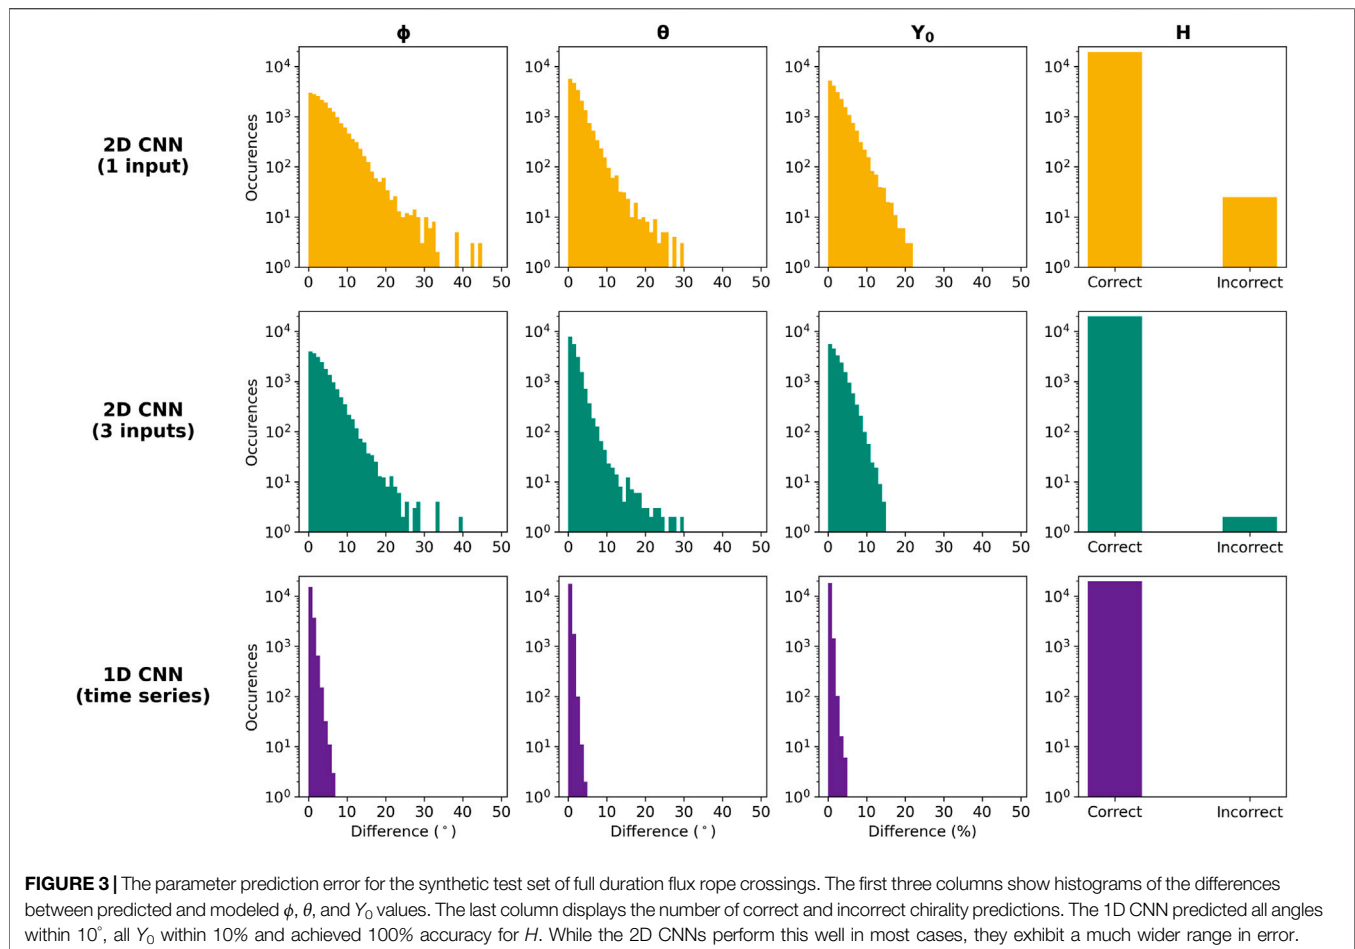
and 20% testing sets. This resulted in 58,800 synthetic flux ropes used for training, 19,600 used for validation, and 19,600 used for testing. The training set was used in a supervised learning fashion with the Adam optimizer (Kingma et al., 2015) with the validation set used during the training process to avoid overfitting. All networks were set to train over 500 epochs, but the 2D CNNs had early stopping from criteria on the validation set at around 35 to 50 epochs. The 1D CNN had a training time of 12 min and both the one input and three input 2D CNNs had training times approaching 4–6 h.

The setup of the second experiment, in which we train over all full and partial flux ropes, was similar. A 60/20/20 split was used, with validation criteria used for early stopping and evaluation on the testing set. Again, the 1D CNN trained over all 500 epochs while the

2D networks reached early stopping within 50 epochs. The 1D CNN took just over 2 h to train, while the 2D CNNs completed in 6–10 h. It should be noted that all percentages of a particular flux rope configuration were included in an input batch. Also, an important consideration in this scenario is that the networks will be seeing multiple inputs that share the same output. All neural networks were constructed, trained, and tested using Python 3.8.10, Keras 2.4.3 (Chollet, 2015), TensorFlow 2.3.1 (Abadi et al., 2015), Numpy 1.18.5 (Harris et al., 2020), and Scipy 1.7.1 (Virtanen et al., 2020).

## 2.3 Wind Spacecraft

The final segment of this work is to evaluate the trained CNNs on flux ropes observed by the Wind spacecraft. This application of the CNNs on non-synthetic data helps us understand the limitations of the flux



rope analytical model and the transition to actual space weather forecasting. Nieves-Chinchilla et al. (2018) carried out a comprehensive study of the internal magnetic field configurations of ICMEs observed by Wind at 1AU in the period 1995-2015. In this analysis, the term magnetic obstacle (MO) is adopted as a more general term than magnetic cloud in describing the magnetic structure embedded in an ICME. The authors used the Magnetic Field Instrument (MFI) (Lepping et al., 1995) and Solar Wind Experiment (SWE) (Ogilvie et al., 1995) to manually set the boundaries of the MO through visual inspection. All MO events were sorted into three broad categories based on the magnetic field rotation pattern: events without evident rotation (E), those with single magnetic field rotation (F), and those with more than one magnetic field rotation (Cx). More recently, Nieves-Chinchilla et al. (2019) presented an in-depth classification, which further classified the F types events into F-, Fr, and F+ based on the angular span of the magnetic field rotation. These events were then manually fit with the Circular-Cylindrical N-C model by a human expert. Of the events cataloged and fit, those that were classified as the Fr type tended to be the ones that could best be fit with the N-C model. Because we restricted our training set of synthetic data to flux rope cases with a  $Y_0 > 0$ , we also restrict our Wind test event cases to this criteria. We use this subset of 75 Wind Fr type events to evaluate our neural network predictions on actual flux rope observations. We compare the human-

fit key parameters to the neural network predictions. While we have high confidence in the human expert's fit values, we acknowledge that they are not definitive. Other experts may parameterize the event slightly differently. Instead of using the human expert as ground-truth, we are interested in seeing if a neural network, trained on the same physical model that the human expert used, will arrive at similar flux rope orientations. The average correlation coefficient is used to compare human and neural network fits to the Wind magnetic field profiles.

As noted earlier, the 1D CNN is configured to input vectors of size 50 and trained on normalized synthetic data, requiring some pre-processing for use with real-event data. We begin with the 1-min resolution MFI data for each of the 75 Wind events and apply a 5-point moving average smoothing followed by interpolation to 50 points evenly spaced in time.

## 3 RESULTS

### 3.1 Full Duration Synthetic Flux Ropes

Results of applying the neural networks trained on full duration flux ropes to the testing set of full duration flux ropes are shown in Figure 3. The  $\phi$ ,  $\theta$ , and  $Y_0$  panels display histograms of the difference between the neural network's predictions and the true values used to create the simulated instance for longitude,

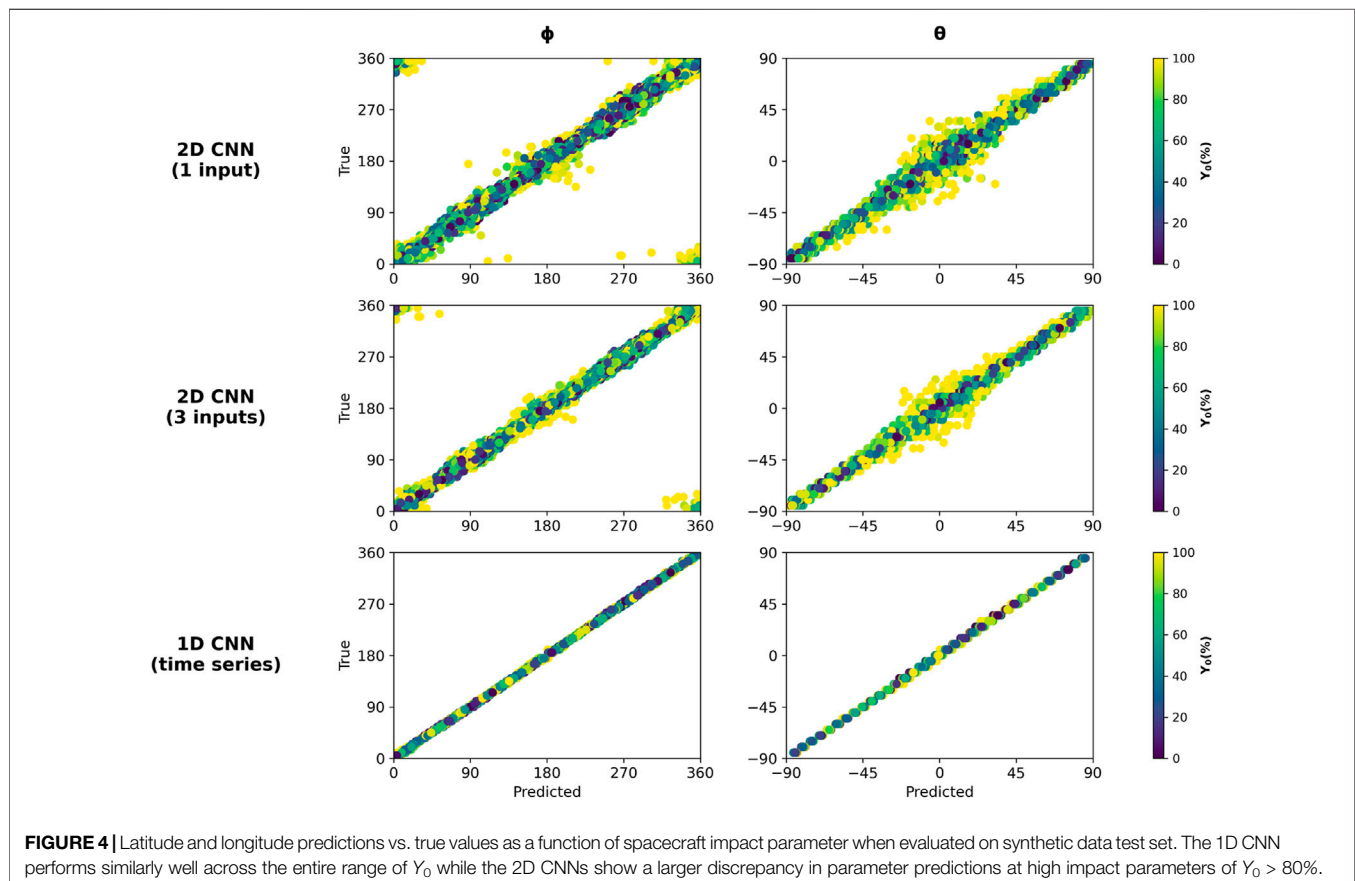
**TABLE 1** | Median of the parameter differences shown in **Figure 3**.

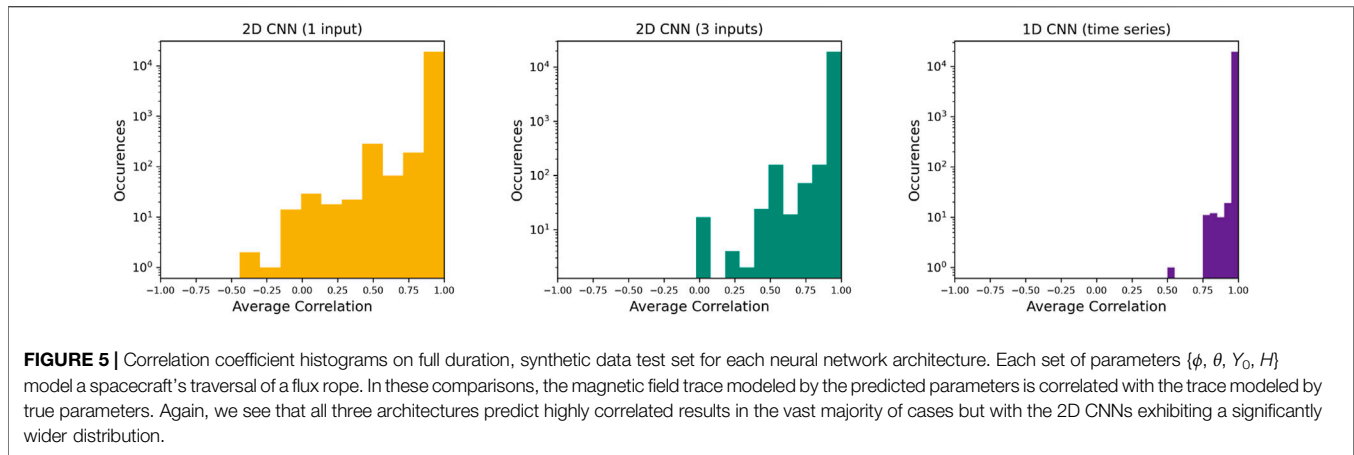
| CNN                | Median difference |        |      |
|--------------------|-------------------|--------|------|
|                    | 2D (1)            | 2D (3) | 1D   |
| $\phi(^{\circ})$   | 3.65              | 2.67   | 0.54 |
| $\theta(^{\circ})$ | 1.86              | 1.31   | 0.37 |
| $Y_0$ (%)          | 2.13              | 1.93   | 0.34 |

latitude, and impact parameter, respectively. The  $H$  panel shows the number of correct and incorrect chirality predictions. Subsequent figures use the same color scheme (2D CNN with 1 input in orange, 2D CNN with 3 inputs in green, and 1D CNN in purple) for clarity. **Table 1** lists the median values of these difference distributions. The 2D CNN with a single input channel has the highest median difference across all three of the re-valued key parameters, as well as the most skewness. The 1D CNN shows the least skewness and the lowest median difference values across the parameters. The 2D CNN with three inputs falls in between, but with median difference and skewness more similar to the other 2D network than the 1D. A similar trend is seen in the  $H$  predictions, with the one input 2D network having the most incorrect classifications and the 1D network making no incorrect classifications. Taken together, it is evident that the 1D CNN, which is applied to the time series directly, gives more accurate predictions across all four output parameters.

While the 1D CNN gives the most accurate predictions, all three architectures give reasonably useful predictions for the vast majority of cases. The bulk of the prediction errors are less than  $15^{\circ}$  for  $\phi$  and  $\theta$  and under 10% for  $Y_0$  for both of the 2D CNNs. **Figure 4** illustrates the prediction errors as a function of  $Y_0$ . The 2D CNNs using hodograms as input have the most significant  $\phi$  and  $\theta$  prediction errors, which occur at large  $Y_0$ . In contrast, the 1D CNN more accurately predicts  $\phi$  and  $\theta$  over the entire range of simulated  $Y_0$ . Clearly, the architecture of the neural network plays a role in prediction accuracy and leads to an important trade off. The two-dimensional networks, by using hodogram input, remove time from the training process. This makes little difference with the synthetic training data but is an advantage when working with data from time-varying, real ICME events, as the data can be used with less manipulation in pre-processing. Yet, this comes at the cost of less accurate predictions at large spacecraft impact parameters ( $Y_0$ ). The trade off is that the simpler and more accurate 1D network comes with the added complexity of determining the most appropriate data transformations to fit the measured time-series to the prescribed input array dimensions of the network.

Our CNNs were each designed with four loss functions and our analysis up to this point has looked at each predicted parameter individually. We now turn our attention to evaluating the predictions as a set. To do so, we use the predicted  $\phi$ ,  $\theta$ ,  $Y_0$ , and  $H$  to reconstruct the magnetic field





time series with the analytical model and correlate it with the simulated magnetic field used as input for the CNN. For analysis, we use the average correlation coefficient,  $r$ , defined as:

$$r = \frac{r_x + r_y + r_z}{3} \quad (1)$$

where  $r_x$  is the Pearson's correlation between the simulated and reconstructed  $b_x$  components,  $r_y$  is the Pearson's correlation between the simulated and reconstructed  $b_y$  components, and  $r_z$  is the Pearson's correlation between the simulated and reconstructed  $b_z$  components.

**Figure 5** shows the average correlation values for each of the three CNN architectures. While the bulk of the simulated data predicted by the 2D CNNs have high average correlation, over 0.75, there is a long tail of predictions with much lower correlation. The single input 2D CNN even makes some predictions that lead to negative correlations. As in the individual key parameters, we see the 1D CNN applied to the time series outperforming the 2D CNNs. In the case of the 1D CNN, we find only one correlation value below 0.75. Further analysis reveals that this event occurred at a simulated  $\phi$  value of  $175^\circ$ . The 1D neural network predicted a  $\phi$  value of  $181^\circ$ . Although the neural network predictions were fairly accurate (within  $5^\circ$ , 5%, and correct  $H$ ), this small deviation in  $\phi$  changed the spacecraft's trajectory through the flux rope leading to a negative correlation in the  $b_z$  component. Because the 2D CNNs are impacted by their difficulties making predictions at large spacecraft impact parameters, we see many of the poor average correlation coefficients in the 2D CNNs at large spacecraft impact parameters.

### 3.2 Partial Duration Synthetic Flux Ropes

In our second experiment, we retrained a second version of each of the three neural networks, this time using the full set of 980,000 full and partial duration flux ropes. Like the difference comparisons shown in **Figure 3** and **Table 1**, **Table 2** provides summary statistics of  $\phi$ ,  $\theta$ , and  $Y_0$  prediction error as a function of percentage of flux rope observed. All three models make fairly accurate predictions even when seeing just 10% of the flux rope and then continue to improve their prediction accuracy

up to a point. After this point, the key parameter accuracy gets worse as higher percentages of the flux ropes are fed to the networks. The level of observation giving the lowest median errors for each CNN is highlighted in yellow, with the next lowest medians highlighted in green. Additionally, all three models were able to predict the correct  $H$  over 99% of the time at all percentages of flux rope observed.

It is worth noting that all three networks perform worse at 100% duration when trained with partial duration flux ropes as compared to these same networks trained only with full duration flux ropes. The introduction of partial flux ropes into the training produces more error (see **Table 1** and **Table 2**). We suspect this is due to multiple inputs now producing the same output. It remains for future research to conduct a more in depth analysis into how to combat this.

As with the networks trained only with full duration flux ropes, the 1D CNN gives better predictions across all parameters. We see a familiar pattern emerge in the 2D CNNs; they have difficulty predicting spacecraft impact parameter and more often predict chirality incorrectly. This in turn leads to greater inaccuracies in  $\phi$  and  $\theta$  predictions. Given that the 1D CNN outperformed the 2D CNNs in both training experiments, we focus only on the 1D architecture when evaluating network performance on actual spacecraft measurements.

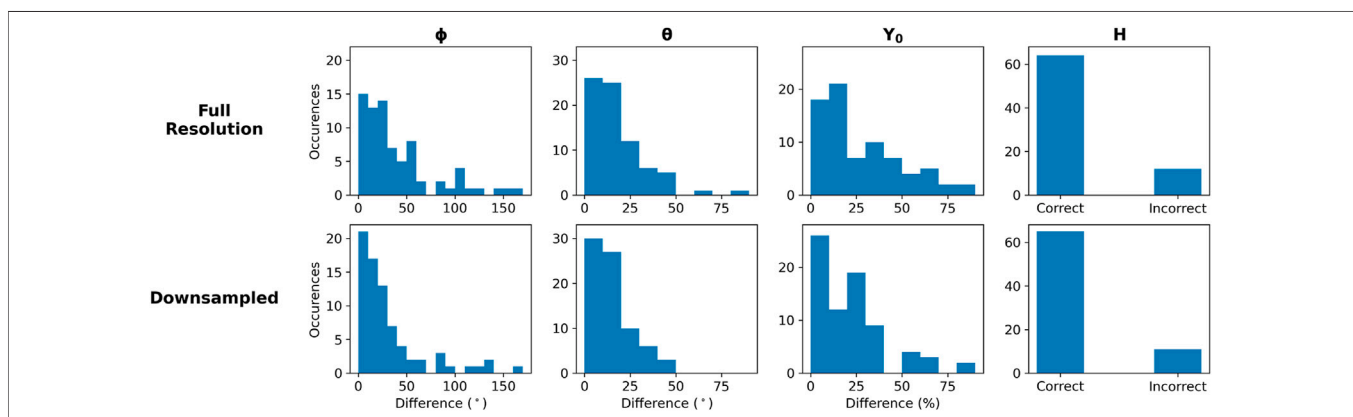
### 3.3 Application to Wind Catalog Flux Ropes

To assess the transfer-ability of this technique to real-time use, we applied the 1D CNN trained on full duration flux ropes to the 75 selected Wind events described in **Section 2.3** with the data processed in two ways. The first approach, which we label Full Resolution, is where we simply use the window smoothing before interpolating the event down to 50 points. The second approach, called Downsampled, first applies 15 min averaging before smoothing and interpolation. The idea being that the Downsampled approach would further reduce fluctuations found inside Wind flux ropes. Comparing Full Resolution and Downsampled would help us isolate the impacts of fluctuations. The difference histograms in **Figure 6** show the result of comparing the fit parameters from Nieves-Chinchilla et al. (2019) (N-C) with the neural network



**TABLE 2 |** Median parameter differences by percentage of flux rope observed for the neural network architectures when trained using partial duration crossings. Cells highlighted in yellow indicate the lowest error for each (CNN, parameter) pair and cells highlighted in green, the next two lowest errors. The overall performance of the 1D CNN continues to be significantly better than the 2D CNNs. The 2D CNNs make their best predictions when seeing less of the flux rope crossing.

| % Observed | 2D (1) |          |       | 2D (3) |          |       | 1D     |          |       |
|------------|--------|----------|-------|--------|----------|-------|--------|----------|-------|
|            | $\phi$ | $\theta$ | $Y_0$ | $\phi$ | $\theta$ | $Y_0$ | $\phi$ | $\theta$ | $Y_0$ |
| 10         | 7.67°  | 5.85°    | 7.33% | 5.49°  | 3.92°    | 5.61% | 2.10°  | 1.23°    | 1.28% |
| 20         | 6.62°  | 4.89°    | 6.62% | 4.94°  | 3.47°    | 5.15% | 1.58°  | 1.02°    | 1.08% |
| 30         | 6.25°  | 4.58°    | 6.20% | 4.70°  | 3.28°    | 4.98% | 1.37°  | 0.90°    | 0.94% |
| 40         | 6.29°  | 4.38°    | 5.99% | 4.60°  | 3.24°    | 4.96% | 1.23°  | 0.84°    | 0.84% |
| 50         | 5.96°  | 4.28°    | 5.87% | 4.61°  | 3.24°    | 4.95% | 1.14°  | 0.81°    | 0.79% |
| 60         | 6.02°  | 4.22°    | 5.75% | 4.70°  | 3.31°    | 4.96% | 1.11°  | 0.80°    | 0.76% |
| 70         | 6.17°  | 4.23°    | 5.79% | 4.74°  | 3.33°    | 5.06% | 1.04°  | 0.76°    | 0.72% |
| 80         | 6.13°  | 4.36°    | 5.89% | 4.83°  | 3.37°    | 5.10% | 1.01°  | 0.75°    | 0.71% |
| 90         | 6.38°  | 4.45°    | 6.13% | 4.93°  | 3.41°    | 5.25% | 1.04°  | 0.76°    | 0.75% |
| 100        | 7.07°  | 4.81°    | 6.52% | 5.17°  | 3.61°    | 5.51% | 1.10°  | 0.79°    | 0.83% |



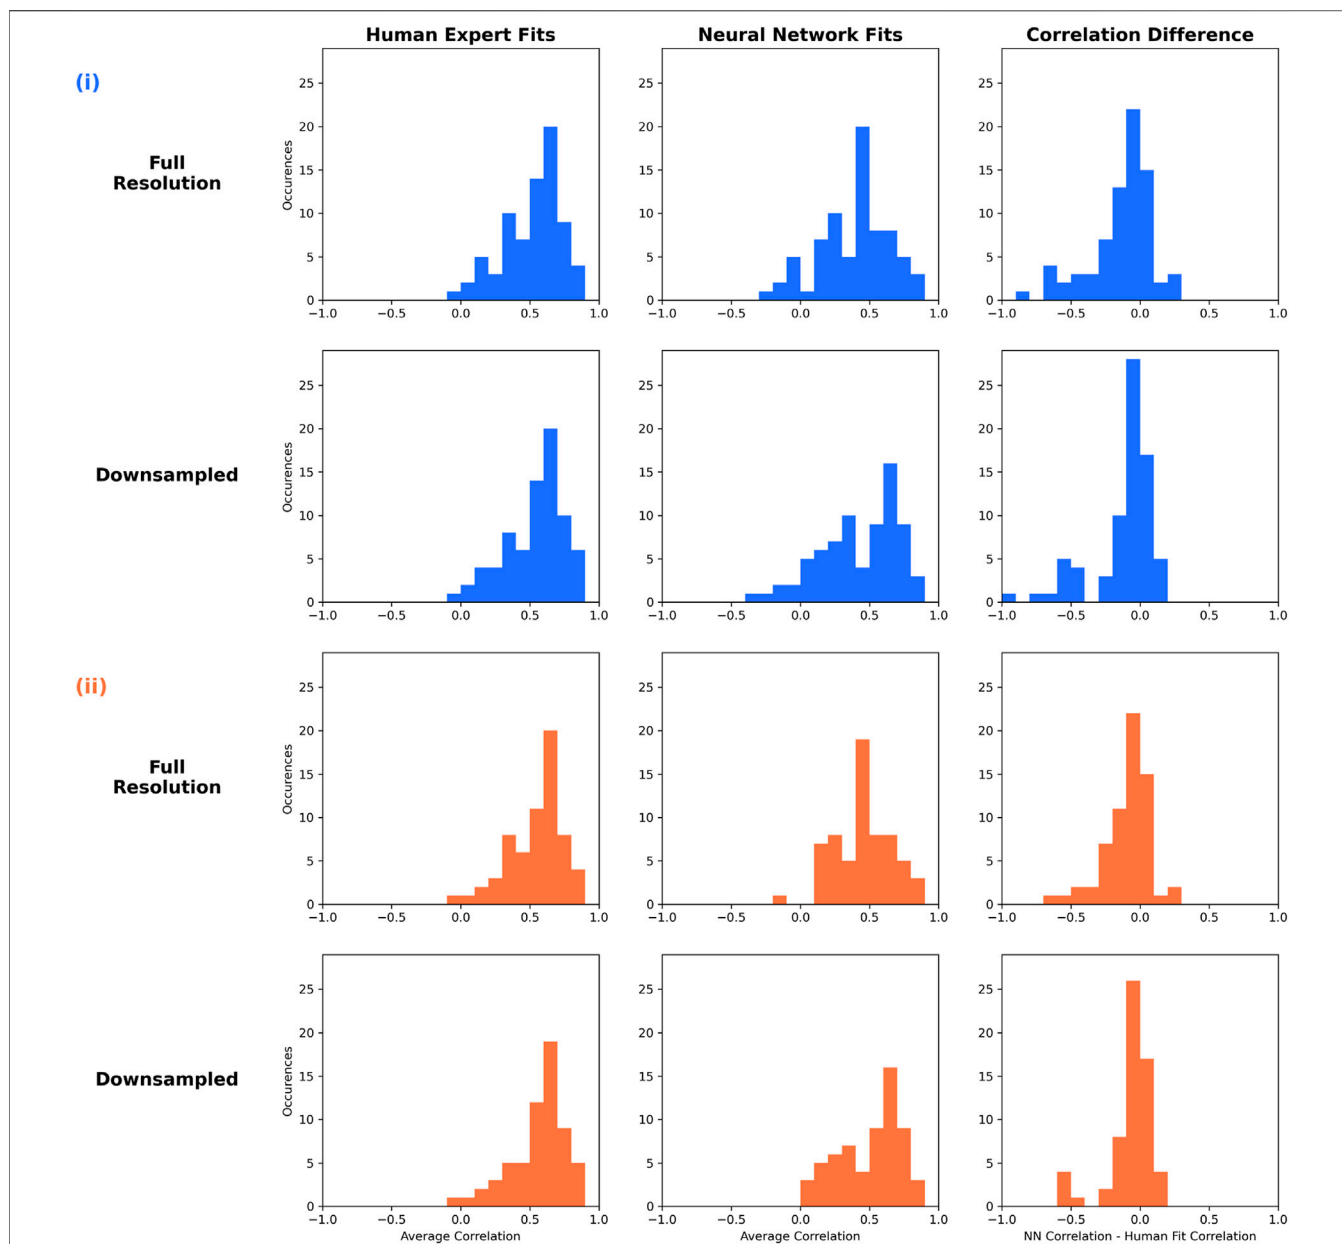
**FIGURE 6 |** The 1D CNN parameter prediction error for the Wind event test set of full duration flux rope crossings. Two sets of predictions were generated: One from processing 1 min Wind MFI measurements and the second from processing the Wind MFI measurements down-sampled to the 15 min averages. The human-fit parameter values from the published ICME catalog were compared against neural network predictions. The overall error magnitude is greater than when tested on synthetic input but shows the same trend. Predictions improve when the down-sampled input is used.

predicted parameters. We note that our neural network was trained with force free synthetic flux ropes ( $C_{10}$  parameter equal to 1). The N-C fittings allowed for deviations from a force free flux rope. This difference likely played a role in the discrepancies between neural network predictions and the human expert’s fits.

Most  $\phi$  predictions were within  $50^\circ$  of the hand-fit value but the maximum error was over  $150^\circ$ . The  $\theta$  errors tend to be less than  $25^\circ$  with a maximum around  $80^\circ$ . Most  $Y_0$  predictions are within 30% of the comparison values with a maximum near 80%. Across all these real-valued key parameters, the predictions made from the Downsampled input display a less skewed error distribution with a higher percentage of the predictions having relatively small error. The network produced similar results predicting chirality ( $H$ ) when fed with Full Resolution and Downsampled input.

We extend this comparison with analysis of average correlation coefficient. We display correlation between the interpolated Wind observations and the magnetic field vectors generated using the N-C fit parameters as well as

the correlation between interpolated Wind observations and magnetic field vectors generated from neural network predictions. **Figure 7** column 1 shows the distribution of average correlation between the human-fit model and Wind data. Column 2 is the distribution of average correlation between the CNN fit and Wind data. Displayed in column 3 is the difference histogram showing the neural network correlation minus the hand-fit correlation for each of the Wind flux rope events. Positive values indicate the neural network produced a statistically more reliable fit. Panel 7(i) shows these distributions for all of the 75 events. Panel 7(ii) shows the distributions when we consider only the events in which the CNN predicted the same chirality as the human-fit. We see good agreement in average correlation coefficients when the predictions are used to reconstruct the magnetic field time series. The shape of the distributions are similar to those from the comparison with human expert fits and an event by event comparison with human expert fits leads to a difference histogram nearly centered at zero. When we look at the Downsampled neural network predictions with



**FIGURE 7** | Correlation distributions and comparisons for the Wind event test set of full duration flux rope crossings. Column 1 shows the correlation score between the human-fit parameters and the Wind measurements. Column 2 shows the correlation between the 1D CNN prediction and the Wind measurements. The third column displays the difference between the human-fit correlations and the CNN prediction correlations. (i) Includes all 75 Wind test cases. (ii) Includes cases where the CNN  $Y_0$  prediction matched human expert’s chirality only. When the CNN predicts  $Y_0$  correctly, the correlation to Wind data is similar to that of the human-expert.

chirality prediction matching the chirality of the human expert (**Figure 7(ii)**) we see no negative correlations.

We next applied the network trained on full and partial duration flux ropes to the aforementioned subset of 75 Wind Fr events. A summary of the results are shown in **Table 3** where we list median differences between network predictions and hand-fit values as a function of flux rope observed. Also shown are the percentage of events where predicted chirality and hand-fit chirality match. The median difference in longitude ranges from  $58^\circ$  to  $89^\circ$ ; in latitude from  $31^\circ$  to  $50^\circ$ ; and in impact

parameter from 36 to 53%. The network predicted the chirality correctly between 52 and 65% of the time.

### 3.4 Number of Wind Events to Train a Network

Experimenting with synthetic and real flux ropes raised an interesting question: How many real flux ropes are needed to train a neural network and how many suitable flux ropes are available for such a study? We can not answer this question

**TABLE 3** | Wind event summary statistics as a function of percentage of flux rope observed for the 1D network trained with both full and partial duration flux ropes. Human-fit parameters are compared to neural network predictions and the  $\phi$ ,  $\theta$ , and  $Y_0$  columns are median differences between the two. The  $H$  column is the percentage of events where the chirality prediction matches hand-fit value.

| % Observed | $\phi$ | $\theta$ | $Y_0$ (%) | $H$ (%) |
|------------|--------|----------|-----------|---------|
| 10         | 89°    | 50°      | 36        | 63      |
| 20         | 66°    | 42°      | 39        | 52      |
| 30         | 69°    | 32°      | 53        | 60      |
| 40         | 64°    | 37°      | 39        | 60      |
| 50         | 70°    | 33°      | 37        | 60      |
| 60         | 69°    | 37°      | 51        | 60      |
| 70         | 73°    | 31°      | 44        | 64      |
| 80         | 73°    | 34°      | 53        | 65      |
| 90         | 73°    | 33°      | 47        | 56      |
| 100        | 58°    | 42°      | 44        | 60      |

conclusively. As discussed in a previous section, and elaborated on below, neural networks trained on synthetic events do not transfer perfectly to Wind. However, we can perform one additional experiment to roughly gauge an answer.

We re-used the train-validation-test split of our synthetic flux ropes mentioned in **Section 2.2**. We then set up a loop of nine iterations. In each iteration, we randomly selected a diminishing subset of the training data, trained a 1D CNN with that subset, and then evaluated the trained model on the testing set of 19,600 synthetic flux ropes. The subsets were selected at random to simulate what happens in practice where we cannot dictate the orientation of flux ropes observed by a spacecraft. The testing consisted of using the trained CNN to make orientation predictions for each of the 19,600 test flux ropes, use those orientation predictions to create the corresponding magnetic field profiles, and correlate those magnetic field profiles with the magnetic field profiles of the test flux rope. As an evaluation metric, we computed the percentage of the test correlations greater than or equal to 0.75. Within each iteration, the subsetting-training-prediction-correlation workflow was repeated three times to investigate how the random subsetting might impact the results.

**Table 4** lists the results. Over ninety percent of testing events have an average correlation coefficient above 0.75 as long as the training set size is over 200 events. Put another way, a 1D CNN trained with roughly 200 events produces average correlation coefficients on par with the 2D 3-input CNN (middle panel of 5). We do note, however, that our experiment is based on training the network with a specific flux rope model and simulated (synthetic) flux ropes. The specific flux rope model chosen will play a role as more complex descriptions of flux ropes (i.e., taking into account compression/expansion) will have more output parameters, which in turn will impact accuracy. In addition, these synthetic flux ropes do not take into account the turbulent fluctuations found in real flux ropes - a further source of prediction error. Nevertheless, it is interesting to note that we may be tantalizingly close to a neural network trained on real observations. There are 151 Wind events in Nieves-Chinchilla et al. (2019) that could potentially be used in training. The HELIO4CAST ICME catalog version 2.1 (Moestl et al., 2020)

**TABLE 4** | Percentage of synthetic flux rope predictions with an average correlation coefficient of 0.75 or greater as a function of training set size. The 1D CNN was used for training. Each training set size was repeated three time, each time taking a different random sample. The percentages reported are the average of the three repetitions. The SD column lists the standard deviation of the three repetitions.

| # Flux ropes in training | % $\geq 0.75$ | SD   |
|--------------------------|---------------|------|
| 29,440                   | 99            | 0.5  |
| 14,592                   | 99            | 0.4  |
| 7,168                    | 98            | 0.05 |
| 3,584                    | 97            | 0.4  |
| 1,792                    | 97            | 1.3  |
| 1,024                    | 97            | 0.6  |
| 512                      | 95            | 0.3  |
| 256                      | 93            | 0.6  |
| 128                      | 88            | 0.88 |

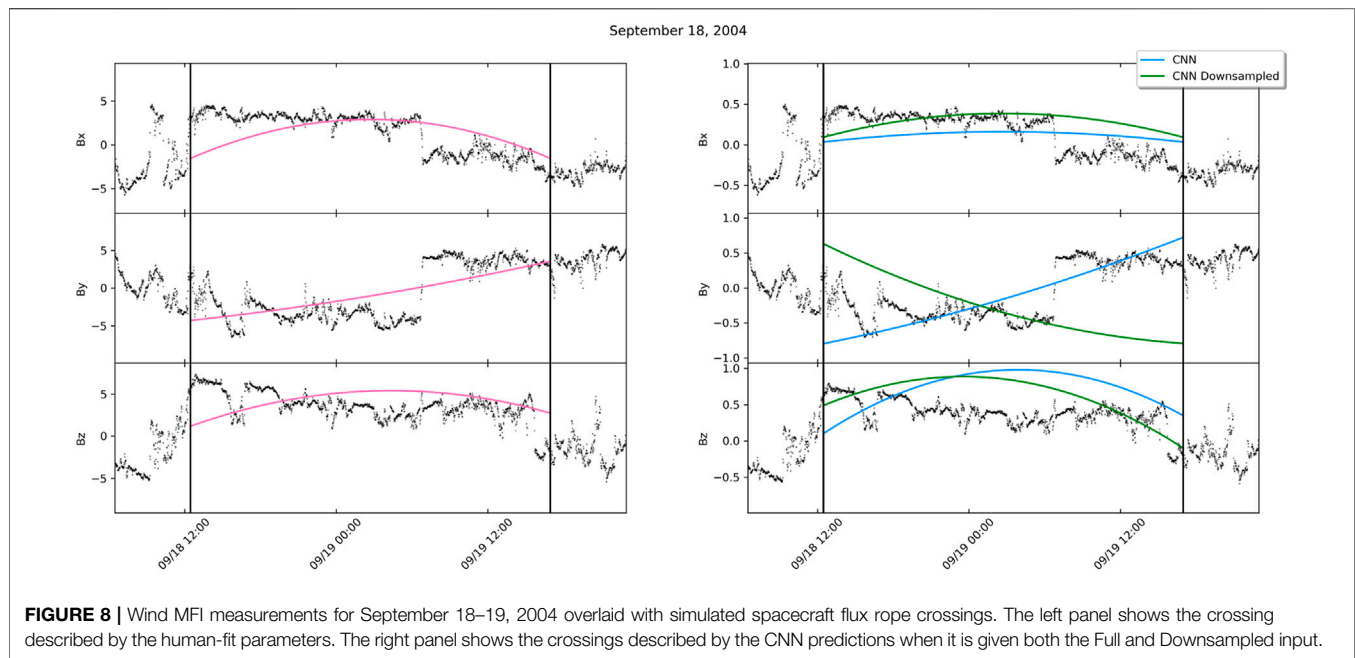
has over 1,000 ICMEs identified from multiple spacecraft. It is unknown how many of these ICMEs have associated flux ropes. Once identified, assuming there are enough, those flux ropes will need to be fit by human experts to provide labeled data for supervised learning. Nevertheless, our experiment provides the intriguing result that a few hundred more events may be all that is needed. Existing ICME catalogs may hold enough events that a concerted effort could lead to training set of real flux ropes in the coming years.

## 4 DISCUSSION AND CONCLUSION

Our experiments have demonstrated that convolutional neural networks are capable of providing extremely reliable characterizations of flux ropes from synthetic data. A trained network can use the structure of simulated magnetic field vectors to learn filters that map to accurate flux rope key parameter predictions; successfully inferring large scale, 3D information from single-point measurements.

When trained only on examples of full duration flux ropes, all three architectures predict key parameters of a flux rope which correlate well with the input data; however the best performing is the 1D network that feeds on time series data. Although the 2D networks that use hodogram style input do not see the same, perfect accuracy in predicting the chirality as the 1D network, the difference is statistically minor. The biggest weakness in the hodogram-input CNNs is when interpreting flux rope traces generated with a high spacecraft impact parameter. It is possible that similarities in hodogram shape profile between low- and high-valued  $Y_0$  are activating similar filters in the 2D networks and leading to poor predictions in these cases.

When we extend the synthetically trained networks to include both partial and full duration traces through flux ropes, we still find this approach highly accurate. The CNNs are capable of making reliable predictions having only seen a fraction of the full flux rope. Although the overall discrepancy between the true and predicted values is higher than when done with only full duration traces, all median differences are well within a tolerable limit. In



these idealized, synthetic, circular-cylindrical flux ropes even the poorest performing network is able to predict orientation angles with a median error under  $8^\circ$  after only observing 10% of a simulated spacecraft crossing. The 1D network here, at only 10% observed, is able to give predictions with lower median difference than the 2D networks do when trained and tested with only full flux ropes. All three models show a peak performance at some point prior to seeing 100% of the flux rope crossing, perhaps due to some similarity in shape between low percentage of observation and high percentage. It is interesting to note that the 2D networks hit their peak predictive point earlier than the 1D CNN, 2D one input at 50–60% and 2D three input even earlier at 40–50%. This suggests that research into where the convolutional network is looking (for example, with the Grad-CAM method (Selvaraju et al., 2017)) can help us further understand the benefits and limitations of hodograms and time series as inputs. Future research will examine where the network is focusing its attention and if this can be exploited for more accurate predictions earlier in the forecasting process.

With the success of the 1D CNN in real-time forecasting from idealized synthetic data, we evaluated this trained 1D network on partial Wind event data. Overall, the neural network struggles to reproduce the accuracy achieved on the synthetic data set. Unlike the synthetic case, we see no trend towards a peak performance point dependent on the amount of flux rope observed. When looked at on a case-by-case basis, there are a few specific events in which the neural network is able to make accurate predictions after only seeing a fraction of the flux rope. In general, however, the median difference in angle and impact parameter prediction falls well outside any tolerance levels for useful prediction and the chirality is only correct approximately 60% of the time. Clearly, the partial-trained CNN cannot be transferred as-is

to real-time application, but insight can be found by examining the results of the full duration network evaluated with Wind events.

Applying the 1D CNN trained only on full duration synthetic flux ropes to *in situ* Wind events, we again see the individual parameter predictions show significant deviation from hand fit values. However, we note lower median differences and higher  $H$  accuracy than when the network trained on both full and partial events was applied to Wind. Using down-sampled input improves this even further. Yet, by looking at the average correlation scores we see that the flux rope analytical model is robust to small deviations - small changes in longitude in particular do not lead to significant differences in reconstructed time series. We also find the neural network robust to variation in solar wind speed, expansion/compression, duration, and to some degree, magnetic field fluctuations. The neural networks were trained on synthetic data that was all generated with a simulated solar wind velocity of 450 km/s and simulated flux rope radius of 0.07 AU; yet, are able to offer reasonable predictions for Wind Fr events having significant differences in solar wind speed, expansion/compression, duration, and magnetic field fluctuations.

The neural network gives reliable predictions in a number of events and exhibits a distribution of average correlations that is qualitatively similar to those from the human expert. As evident in the right-most column of **Figure 7**, the neural network results in better average correlation in nearly half of the 75 events. When we consider only cases in which the network prediction for  $H$  matches the human-fit  $H$  the correlation to Wind data is even greater.

Analysis reveals two primary reasons the neural network performs less accurately on Wind events; incorrect physical model (Wind flux ropes not fitting the circular cylindrical

assumptions) and internal physical processes (such as fluctuations and discontinuities) that alter the expected magnetic field profile of a smooth flux rope. An example of a flux rope with magnetic field fluctuations and a discontinuity is shown in the event with MO beginning on 18 September 2004 in **Figure 8**. Down-sampling the Wind magnetic field data from 1-min to 15-min prior to interpolating to 50 points reduces the difference between neural network predictions and hand-fit values. The down-sampling further smooths out the magnetic field time series removing small-scale fluctuations. However, down-sampling cannot account for all observed internal physical processes that lead to a deviation from the expected smooth flux rope profile. The September 2004 event illustrates how differences in data processing can have a strong effect on the resulting prediction. In this particular example, the predictions made from the Wind data without prior averaging match the hand fit predictions well, while those from the down-sampled input clearly lost important information. The choice of 15-min averaging was arbitrary and is presented here to highlight how data pre-processing can have both positive and negative impacts on prediction accuracy. It remains for future research to systematically address fluctuations and determine an optimal input resolution.

Of the total 151 Wind Fr events in Nieves-Chinchilla et al. (2019), only 41% were classified as a flux rope by the neural network developed in dos Santos et al. (2020) when trained with no fluctuations. This same network classified 84 and 76% as flux ropes when trained with synthetic data augmented with 5 and 10% Gaussian fluctuations, respectively. In other words, some of the Wind events on which we do poorly finding good parametrization, would not have been considered a flux rope by the first step of an automated fitting workflow. At present, magnetic field fluctuations are not fully accounted for in flux rope analytical models and pre-processing of neural network input data does not fully address the discrepancy between synthetic and spacecraft observed flux ropes. Accurately accounting for fluctuations in measured data appears to be a significant factor for improving an automated space weather forecasting pipeline. Early experimentation with 5% Gaussian fluctuations in our study did not lead to significant improvement. Solar wind and flux rope turbulence is known to be non-Gaussian. Yet, at present, a complete understanding of turbulence leads analytical models lacking in this regard. We choose to not introduce non-realistic fluctuations and instead will explore physics-based turbulence enhancements to the analytical model in future research.

The ultimate source of prediction error in any CNN is in the inputs not matching any of the learned filters. In the case of Wind events, we notice that the neural network trained with only full duration flux ropes incorrectly predicts chirality in nearly 20% of Wind events. This leads to poor correlation coefficients as the reconstructed time series do not match the Wind observations. Yet, across all implementations of the CNNs with synthetic data the CNNs overwhelmingly identify the correct chirality. This indicates that the convolutional filters the network learned to predict chirality do not transfer to Wind events; that the filters

learned to focus on a quality in the synthetic data that is not shared in the real observations. Interestingly, down-sampling has no effect on chirality predictions. We believe this source of error is related to the physical model chosen to simulate the flux ropes. Wind flux ropes show deviations from the circular cylindrical assumption. This opens the door to tantalizing future evaluations of physics-based flux rope models using an ensemble of neural networks, each trained with a different physical model.

Partial duration predictions and real-time forecasting are not really feasible at this time due in large part to features in the real data that are not present in the training set, though the concept of using CNNs to infer 3D geometric parameters from an *in situ* measurement have been borne out. Additionally, the neural networks have helped highlight the limitations of the physics-based model and even suggested better fittings of some Wind flux ropes. Future work will include implementing a single, physics-based loss function into the CNN to replace the four separate loss functions in the current design as well as enhancing analytical flux rope models to produce training data that includes more realistic turbulence and asymmetry.

## DATA AVAILABILITY STATEMENT

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

## AUTHOR CONTRIBUTIONS

The authors confirm contribution to the paper as follows: study conception and design: TN, AN, LS, and TN-C; flux rope analytical model development and coding: TN-C developed the analytical model and initial version of the code. AN ported the code to language/version used in this study; synthetic data generation: TN; neural network design and implementation; TN, AN, LS; neural network training and testing: TN; analysis and interpretation of results: TN, AN, LS, and TN-C; draft manuscript preparation: TN, AN, LS, and TN-C. All authors reviewed the results and approved the final version of the manuscript.

## FUNDING

LS was supported by NASA Grant 80NSSC20K1580.

## ACKNOWLEDGMENTS

The first author would like to acknowledge Ron Lepping, Daniel Berdichevsky, and Chin-Chun Wu for their many helpful discussions surrounding flux rope physics during earlier projects involving flux rope detection and analysis.

## REFERENCES

- Baker, D. N., and Lanzerotti, L. J. (2008). A Continuous L1 Presence Required for Space Weather. *Space Weather* 6, 1. doi:10.1029/2008SW000445
- Burlaga, L. F. (1988). Period Doubling in the Outer Heliosphere. *J. Geophys. Res.* 93, 4103–4106. doi:10.1029/JA093iA05p04103
- Burlaga, L., Sittler, E., Mariani, F., and Schwenn, R. (1981). Magnetic Loop Behind an Interplanetary Shock: Voyager, Helios, and Imp 8 Observations. *J. Geophys. Res.* 86, 6673–6684. doi:10.1029/JA086iA08p06673
- Camporeale, E. (2019). The Challenge of Machine Learning in Space Weather: Nowcasting and Forecasting. *Space Weather* 17, 1166–1207. doi:10.1029/2018SW002061
- [Software] Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., et al. (2015). *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. Accessed on: January 2021.
- [Software] Chollet, F. (2015). Keras. Available at: <https://keras.io> (Accessed on: January 2021).
- dos Santos, L. F. G., Narock, A., Nieves-Chinchilla, T., Nuñez, M., and Kirk, M. (2020). Identifying Flux Rope Signatures Using a Deep Neural Network. *Sol. Phys.* 295, 131. doi:10.1007/s11207-020-01697-x
- Gosling, J. T., Pizzo, V., and Bame, S. J. (1973). Anomalous Low Proton Temperatures in the Solar Wind Following Interplanetary Shock Waves—Evidence for Magnetic Bottles? *J. Geophys. Res.* 78, 2001–2009. doi:10.1029/JA078i013p02001
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., et al. (2020). Array Programming with NumPy. *Nature* 585, 357–362. doi:10.1038/s41586-020-2649-2
- Jian, L., Russell, C. T., Luhmann, J. G., and Skoug, R. M. (2006). Properties of Interplanetary Coronal Mass Ejections at One AU during 1995 - 2004. *Sol. Phys.* 239, 393–436. doi:10.1007/s11207-006-0133-2
- Kilpua, E., Koskinen, H. E. J., and Pulkkinen, T. I. (2017). Coronal Mass Ejections and Their Sheath Regions in Interplanetary Space. *Living Rev. Sol. Phys.* 14, 5–83. doi:10.1007/s41116-017-0009-6
- Kilpua, E., Koskinen, H. E. J., and Pulkkinen, T. I. (2017). Coronal Mass Ejections and Their Sheath Regions in Interplanetary Space. *Living Rev. Sol. Phys.* 14, 33. doi:10.1007/s41116-017-0009-6
- Kingma, D., and Ba, J. (2015). “Adam: A Method for Stochastic Optimization,” in 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015. Editors Y. Bengio and Y. LeCun (Conference Track Proceedings).
- Klein, L. W., and Burlaga, L. F. (1982). Interplanetary Magnetic Clouds at 1 au. *J. Geophys. Res.* 87, 613–624. doi:10.1029/JA087iA02p00613
- LeCun, Y., and Bengio, Y. (1995). “Convolutional Networks for Images, Speech, and Time Series,” in *The Handbook of Brain Theory and Neural Networks*. Cambridge, MA: MIT Press, 3361, 1995.
- Lepping, R. P., Acuña, M. H., Burlaga, L. F., Farrell, W. M., Slavin, J. A., Schatten, K. H., et al. (1995). The Wind Magnetic Field Investigation. *Space Sci. Rev.* 71, 207–229. doi:10.1007/BF00751330
- Lepping, R. P., Jones, J. A., and Burlaga, L. F. (1990). Magnetic Field Structure of Interplanetary Magnetic Clouds at 1 AU. *J. Geophys. Res.* 95, 11957–11965. doi:10.1029/JA095iA08p11957
- Manchester, W., Kilpua, E., Liu, Y., Lugaz, N., Riley, P., Torok, P. T., et al. (2017). The Physical Processes of Cme/icme Evolution. *Space Sci. Rev.* 212 (3–4), 1159. doi:10.1007/s11214-017-0394-0
- Moestl, C., Weiss, A., Bailey, R., and Reiss, M. (2020). Helio4cast Interplanetary Coronal Mass Ejection Catalog v2.1. figshare. Dataset. doi:10.6084/m9.figshare.6356420.v11
- Nguyen, G., Fontaine, D., Aunai, N., Vandenbosche, J., Jeandet, A., Lemaitre, G., et al. (2018). “Machine Learning Methods to Identify Icmes Automatically,” in 20th EGU General Assembly, EGU2018, Proceedings from the Conference Held, Vienna, Austria, 4–13 April, 2018, 1963.
- Nieves-Chinchilla, T., Jian, L. K., Balmaceda, L., Vourlidas, A., dos Santos, L. F. G., and Szabo, A. (2019). Unraveling the Internal Magnetic Field Structure of the Earth-Directed Interplanetary Coronal Mass Ejections during 1995 - 2015. *Sol. Phys.* 294, 89. doi:10.1007/s11207-019-1477-8
- Nieves-Chinchilla, T., Linton, M. G., Hidalgo, M. A., Vourlidas, A., Savani, N. P., Szabo, A., et al. (2016). A Circular-Cylindrical Flux-Rope Analytical Model for Magnetic Clouds. *Astrophysical J.* 823, 27. doi:10.3847/0004-637X/823/1/27
- Nieves-Chinchilla, T., Vourlidas, A., Raymond, J. C., Linton, M. G., Al-haddad, N., Savani, N. P., et al. (2018). Understanding the Internal Magnetic Field Configurations of ICMs Using More Than 20 Years of Wind Observations. *Sol. Phys.* 293, 25. doi:10.1007/s11207-018-1247-z
- Ogilvie, K. W., Chornay, D. J., Fritzenreiter, R. J., Hunsaker, F., Keller, J., Lobell, J., et al. (1995). SWE, A Comprehensive Plasma Instrument for the Wind Spacecraft. *Space Sci. Rev.* 71, 55–77. doi:10.1007/BF00751326
- Reiss, M. A., Möstl, C., Bailey, R. L., Rüdiger, H. T., Amerstorfer, U. V., Amerstorfer, T., et al. (2021). Machine Learning for Predicting the Bz Magnetic Field Component from Upstream *In Situ* Observations of Solar Coronal Mass Ejections. *Space Weather* 19. doi:10.1029/2021SW002859
- Rodríguez-García, L., Gómez-Herrero, R., Zouganelis, I., Balmaceda, L., Nieves-Chinchilla, T., Dresing, N., et al. (2021). The Unusual Widespread Solar Energetic Particle Event on 2013 August 19: Solar Origin and Particle Longitudinal Distribution. *Astron. Astrophysics* 653, A137. doi:10.1051/0004-6361/202039960
- Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., and Batra, D. (2017). “Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization,” in 2017 IEEE International Conference on Computer Vision (ICCV), Venice, Italy, October 22–29, 2017, 618–626. doi:10.1109/ICCV.2017.74
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., et al. (2020). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat. Methods* 17, 261–272. doi:10.1038/s41592-019-0686-2
- Vourlidas, A. (2014). The Flux Rope Nature of Coronal Mass Ejections. *Plasma Phys. Control Fusion* 56, 064001. doi:10.1088/0741-3335/56/6/064001

**Conflict of Interest:** Author AN is employed by ADNET Systems Inc. and is contracted to NASA/Goddard Space Flight Center where a portion of her employment is to carry out scientific software development and data analysis support. Despite the affiliation with a commercial entity, author AN has no relationships that could be construed as a conflict of interest related to this work.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

**Publisher’s Note:** All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2022 Narock, Narock, Dos Santos and Nieves-Chinchilla. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.