



## RESEARCH ARTICLE

10.1029/2022SW003189

### Key Points:

- Machine learning is used to develop models from unique density data sets
- We compared model predictions along the CHALLENGING Minisatellite Payload (CHAMP) orbit during the 2003 Halloween storms
- We find that models developed on CHAMP and High Accuracy Satellite Drag Model density data can capture density depletion in the post-storm period

### Correspondence to:

R. J. Licata,  
rjlicata@mix.wvu.edu

### Citation:

Licata, R. J., Mehta, P. M., Weimer, D. R., Drob, D. P., Tobiska, W. K., & Yoshii, J. (2022). Science through machine learning: Quantification of post-storm thermospheric cooling. *Space Weather*, 20, e2022SW003189. <https://doi.org/10.1029/2022SW003189>

Received 12 JUN 2022  
Accepted 11 SEP 2022

# Science Through Machine Learning: Quantification of Post-Storm Thermospheric Cooling

Richard J. Licata<sup>1</sup> , Piyush M. Mehta<sup>1</sup> , Daniel R. Weimer<sup>2</sup> , Douglas P. Drob<sup>3</sup> ,  
W. Kent Tobiska<sup>4</sup> , and Jean Yoshii<sup>4</sup>

<sup>1</sup>Department of Mechanical and Aerospace Engineering, West Virginia University, Morgantown, WV, USA, <sup>2</sup>Center for Space Science and Engineering Research, Virginia Tech, Blacksburg, VA, USA, <sup>3</sup>Space Science Division, US Naval Research Laboratory, Washington, DC, USA, <sup>4</sup>Space Environment Technologies, Pacific Palisades, CA, USA

**Abstract** Machine learning (ML) models are universal function approximators and—if used correctly—can summarize the information content of observational data sets in a functional form for scientific and engineering applications. A benefit to ML over parametric models is that there are no a priori assumptions about particular basis functions which can potentially limit the phenomena that can be modeled. In this work, we develop ML models on three data sets: the Space Environment Technologies High Accuracy Satellite Drag Model (HASDM) density database, a spatiotemporally matched data set of outputs from the Jacchia-Bowman 2008 Empirical Thermospheric Density Model (JB2008), and an accelerometer-derived density data set from CHALLENGING Minisatellite Payload (CHAMP). These ML models are compared to the Naval Research Laboratory Mass Spectrometer and Incoherent Scatter radar (NRLMSIS 2.0) model to study the presence of post-storm cooling in the middle-thermosphere. We find that both NRLMSIS 2.0 and JB2008-ML do not account for post-storm cooling and consequently perform poorly in periods following strong geomagnetic storms (e.g., the 2003 Halloween storms). Conversely, HASDM-ML and CHAMP-ML do show evidence of post-storm cooling indicating that this phenomenon is present in the original data sets. Results show that density reductions up to 40% can occur 1–3 days post-storm depending on the location and strength of the storm.

**Plain Language Summary** Machine learning (ML) models are valuable universal function approximators and—if used correctly—can provide scientific information related to the data set used for fitting. A benefit to ML over other common models is that there are no background functions limiting what phenomena it can represent. In this work, we develop ML models on three data sets: the Space Environment Technologies High Accuracy Satellite Drag Model (HASDM) density database, a spatiotemporally matched data set of outputs from the Jacchia-Bowman 2008 Empirical Thermospheric Density Model (JB2008), and an accelerometer-derived density data set from CHALLENGING Minisatellite Payload (CHAMP). These ML models are compared to the Naval Research Laboratory Mass Spectrometer and Incoherent Scatter radar (NRLMSIS 2.0) model to study the presence of post-storm cooling in the upper atmosphere. We find that both NRLMSIS 2.0 and JB2008-ML do not account for post-storm cooling and perform poorly in periods following strong geomagnetic storms. Conversely, HASDM-ML and CHAMP-ML do show evidence of post-storm cooling indicating that this phenomenon is present in the original data sets.

## 1. Introduction

Machine learning (ML) has become increasingly prevalent across many scientific domains. This has been aided largely by improvements in computing power and the use of Graphical Processing Units (GPUs) in model training (Chellapilla et al., 2006). In the space weather community, ML has been used to develop models for problems such as solar flare prediction (Florios et al., 2018), ionospheric scintillation detection (Jiao et al., 2017), and geomagnetic index forecasts (Wing et al., 2005). However, its use is often limited to problem solving, not for investigative purposes. Further, there are rarely any studies on what the model has learned outside of determining its performance metrics. Convolutional neural networks—specifically related to image processing—are inherently easier to understand, as the filters (or weights) can be displayed as images and therefore interpreted (Krizhevsky et al., 2012). This is not a luxury associated with ML regression models where inputs do not have visual qualities, so the weights are difficult to interpret. This motivates this work as it pertains to space weather and the thermosphere.

© 2022. The Authors.

This is an open access article under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

The thermosphere is the neutral region of the upper atmosphere consisting of many atomic and molecular species. Their relative abundance and contribution to the total mass density can vary as a function of altitude, solar activity, and space weather conditions (Emmert, 2015). During abrupt space weather events like geomagnetic storms, the local distributions of the constituents can vary considerably causing significant changes in the total mass density (Fuller-Rowell et al., 1997). Zesta and Oliveira (2019) found that when storms become stronger, the thermosphere both heats and cools at a faster rate. Significant research has been done into a potential cause of the cooling effects, overproduction of nitric oxide (NO) and its infrared emissions.

Kockarts (1980) investigated the cooling impact of the thermosphere due to downward heat conduction, atomic oxygen (O), and NO during a geomagnetic storm in 1974. They found that the reduction in thermopause temperature from the introduction of NO cooling was 440 K, while the addition of O cooling only reduced the temperature by another 35 K. This topic has gained much more attention in recent years due to the NO emission data from the Sounding of the Atmosphere using Broadband Emission Radiometry (SABER) instrument (Russell et al., 1999) and high fidelity density estimates from satellite such as CHALLENGING Minisatellite Payload (CHAMP) (Reigber et al., 2002) and Gravity Recovery and Climate Experiment (GRACE) (Bettadpur, 2012).

Mlynczak et al. (2003) used SABER data during the storm periods of April 2002 and found that NO emissions were notably enhanced during this period. Lei et al. (2012) considered the prominent 2003 Halloween storms to provide a comparison of SABER data to density estimates from both CHAMP and GRACE. They noted a 23%–26% maximum density depletion during the recovery phase for the satellites relative to quiet pre-storm values, and the NO cooling rates during this period remained at a high level. Knipp et al. (2017) examined 192 geomagnetic events to compare NO and neutral density data from GRACE. Their data-based study suggests shock-led interplanetary coronal mass ejections result in an overproduction of NO which provides a cooling effect that compensates for the strong thermospheric expansion that occurs during these storms. The driving force behind the cooling effect is still an active area of research and other mechanisms (e.g., ionosphere-driven atomic oxygen reductions) have been proposed to explain the phenomena (Lei et al., 2021; Mikhailov & Perrone, 2020). We do not attempt to confirm any driving mechanisms in this work.

Using ML, we can investigate the *presence* of post-storm cooling in various data sets and which model drivers may be required to capture it. We first explain the data and models used for model development and comparison. Then, we describe the ML model development process and how we use them to examine this phenomena. We show model predictions during a prominent geomagnetic storm to motivate the importance of this work and provide a quantitative analysis on the effect of geomagnetic time history on the predicted density.

## 2. Data, Models, and Methods

### 2.1. Data and Models

As a benchmark, we used the Naval Research Laboratory Mass Spectrometer and Incoherent Scatter radar (NRLMSIS 2.0) empirical thermosphere model (Emmert et al., 2021), the most recent version of MSIS models dating back to the original MSIS-86 model (Hedin, 1987). NRLMSIS 2.0 was developed as an improvement over the previous MSIS model (NRLMSISE-00). Recent satellite data for temperature, atomic oxygen, and atomic hydrogen was incorporated from the lower atmosphere. The most significant change to the formulation is the coupling between density in the lower atmosphere and thermosphere, dependent on temperature.

NRLMSIS 2.0 uses the *ap* index to account for geomagnetic activity. The *ap* index is indicative of global geomagnetic activity and has a 3 hr cadence. There are two *ap* options when running NRLMSIS 2.0: use only the daily average (known as *Ap*) and current 3-hr value or use a time history of the index. This time history includes *Ap*, current *ap*, *ap*<sub>3</sub>, *ap*<sub>6</sub>, *ap*<sub>9</sub>, *ap*<sub>12–33</sub>, and *ap*<sub>36–57</sub>. The single numerical subscripts refer to the value of the index that many hours prior to the epoch. The combination of two numbers in the subscript refers to the average value over that many hours prior to the epoch (e.g., *ap*<sub>12–33</sub> is the average *ap* value from 12 to 33 hr prior to the epoch). This nomenclature for geomagnetic drivers will be used throughout this manuscript.

We also developed machine-learned models (to be described in Section 2.2) based on three data sets. The first is outputs of the Jacchia-Bowman 2008 Empirical Thermospheric Density Model (JB2008) from the start of 2000 to the end of 2019 (Bowman et al., 2008). We evaluated JB2008 every three hours and at a fixed grid of 12,312 locations including altitude. The longitude, latitude, and altitude resolutions are 15°, 10°, and 25 km, respectively,

with the altitude ranging from 175 to 825 km. JB2008 is driven by four solar indices/proxies. As with NRLMSIS 2.0, it uses  $F_{10.7}$  as a driver for solar activity.  $F_{10.7}$  is a valuable solar proxy that represents the 10.7 cm solar radio emission (Covington, 1948). JB2008 also uses  $S_{10.7}$ ,  $M_{10.7}$ , and  $Y_{10.7}$ , which are explained by Tobiska et al. (2008) and Bowman et al. (2008). We do not go into their details in this work as it is not pertinent to this analysis. For geomagnetic activity, JB2008 uses both  $ap$  and  $Dst$ .  $Dst$  is an index for the strength of the ring current which serves as a useful proxy for geomagnetic activity.

The second data set is the Space Environment Technologies (SET) High Accuracy Satellite Drag Model (HASDM) density database, the first major release of outputs from the US Air Force's HASDM model (Tobiska et al., 2021). This HASDM data set is spatiotemporally matched to the JB2008 outputs (same time period, time resolution, and locations). HASDM is often considered the state-of-the-art for thermospheric density modeling as it assimilates observed drag data from calibration satellites to make corrections to its background empirical model, JB2008 (Storz et al., 2005).

The final data set is the Mehta et al. (2017) accelerometer-derived density estimates from CHAMP. CHAMP was in orbit from 2000 to 2010 with a high inclination and altitude range of 300–460 km. Unlike JB2008 and HASDM, this is an in-situ data set with a much higher cadence—10 s. There are numerous other data sets that have been developed using CHAMP accelerometer data (Bruinsma & Biancale, 2003; Doornbos, 2012; Liu et al., 2005; March et al., 2019; Sutton, 2008), but we proceed with the described data set due to its use in previous work (Licata et al., 2021).

## 2.2. Machine Learning

As the goal is to study the effects of geomagnetic time histories, we do not need to develop a surrogate model for NRLMSIS 2.0 (see Section 2.1). We do, however, proceed with model development for the JB2008, HASDM, and CHAMP data sets. The process for JB2008 and HASDM is identical as they have the same time and space resolution. In addition, HASDM is rooted in JB2008, so we use the same inputs for both data sets. To keep model size reasonable, we leverage principal component analysis (PCA) to reduce the spatial dimensionality from 12,312 to 10 PCA coefficients. This has been demonstrated on similar data sets in the past for the development of thermospheric reduced-order models (Licata et al., 2022; Mehta & Linares, 2017; Mehta et al., 2018). For information on the data preparation and PCA, the reader is referred to Licata et al. (2021) where it was used to analyze the HASDM database. The major difference for CHAMP is that it is an in-situ data set, so location is now an input as opposed to being embedded in the model output.

We first prepared the data for ML, defining the input and output data structures. JB2008 and HASDM have 26 inputs, defined in Table 1. Eight of the inputs are the four solar drivers described in Section 2.1 along with their 81-day centered averages—marked with an “81c” subscript. These inputs are also shared with the CHAMP model. For geomagnetic activity, both JB2008 and HASDM use time histories for  $ap$  and  $Dst$ . The  $ap$  time series is the one described for NRLMSIS 2.0, and the  $Dst$  time series was determined in previous work (Licata et al., 2022). Since CHAMP has a much higher time resolution (1 min), we forgo the use of these geomagnetic indices and instead use  $SYM-H$ , which represents the longitudinally symmetric geomagnetic field disturbances (Iyemori, 1990; Siciliano et al., 2021). The time history was chosen to be similar to those used by the other models. We also use the Poynting flux totals in the northern and southern hemispheres ( $S_N$  and  $S_S$ ) generated by the W05 electrodynamics model (Weimer, 2005a, 2005b). The time and location inputs are described in Equations 1 and 2. The general form of  $(2\pi x/y)$  allows for the linearly increasing inputs to be continuous about their boundaries.  $t_1$  and  $t_2$  represent annual variations, using the day of year (doy), and  $t_3$  and  $t_4$  represent diurnal variations, using universal time (UT). Equation 2 represents the local solar time inputs.

$$t_1 = \sin\left(\frac{2\pi \text{doy}}{365.25}\right), \quad t_2 = \cos\left(\frac{2\pi \text{doy}}{365.25}\right), \quad t_3 = \sin\left(\frac{2\pi \text{UT}}{24}\right), \quad t_4 = \cos\left(\frac{2\pi \text{UT}}{24}\right). \quad (1)$$

$$LST_1 = \sin\left(\frac{2\pi LST}{24}\right), \quad LST_2 = \cos\left(\frac{2\pi LST}{24}\right) \quad (2)$$

The outputs for JB2008 and HASDM are their respective 10 PCA coefficients, while the CHAMP outputs are the local density estimates. With the inputs and outputs set up in an ML format, we can determine an architecture for each data set using Keras Tuner, an automated tool for hyperparameter searching (O'Malley et al., 2019). The

**Table 1**  
*List of Inputs for the Machine Learning Models*

JB2008/HASDM		
Solar	Geomagnetic	Temporal
$F_{10}, S_{10},$	$a_{pA}, a_p, a_{p3},$	$t_1, t_2$
$M_{10}, Y_{10},$	$a_{p6}, a_{p9}, a_{p12-33},$	$t_3, t_4$
$F_{81c}, S_{81c},$	$a_{p36-57}, Dst_A, Dst,$	
$M_{81c}, Y_{81c}$	$Dst_3, Dst_6, Dst_9,$	
	$Dst_{12}, Dst_{15}, Dst_{18}, Dst_{21}$	
CHAMP		
Solar	Geomagnetic	Spatial/Temporal
$F_{10}, S_{10},$	$SYM-H, SYM-H_{0-3}$	$LST_1, LST_2,$
$M_{10}, Y_{10},$	$SYM-H_{3-6}, SYM-H_{6-9}$	$LAT, ALT,$
$F_{81c}, S_{81c},$	$SYM-H_{9-12}, SYM-H_{12-33}$	$t_1, t_2,$
$M_{81c}, Y_{81c}$	$SYM-H_{33-57}, S_N, S_S$	$t_3, t_4$

Note. Note: *LAT* and *ALT* are the latitude and altitude at epoch, respectively.

tuner is provided with a hyperparameter space (e.g., choices for the number of layers, neurons, activation functions, and optimizers) and trains three models with random weight initialization for 100 architectures or trials. The first 25 architectures are randomly selected from the search space. After the random search, the tuner chooses all subsequent architectures using a Bayesian optimization scheme attempting to minimize the loss on the validation set. The 10 best models are evaluated on the training and validation sets to determine the final model considering the lowest errors. The model development for CHAMP, HASDM, and subsequently JB2008 is outlined by Licata and Mehta (2022) with the only difference being the *SYM-H* time series inputs for CHAMP. We provide additional details on model development in Appendices A and B.

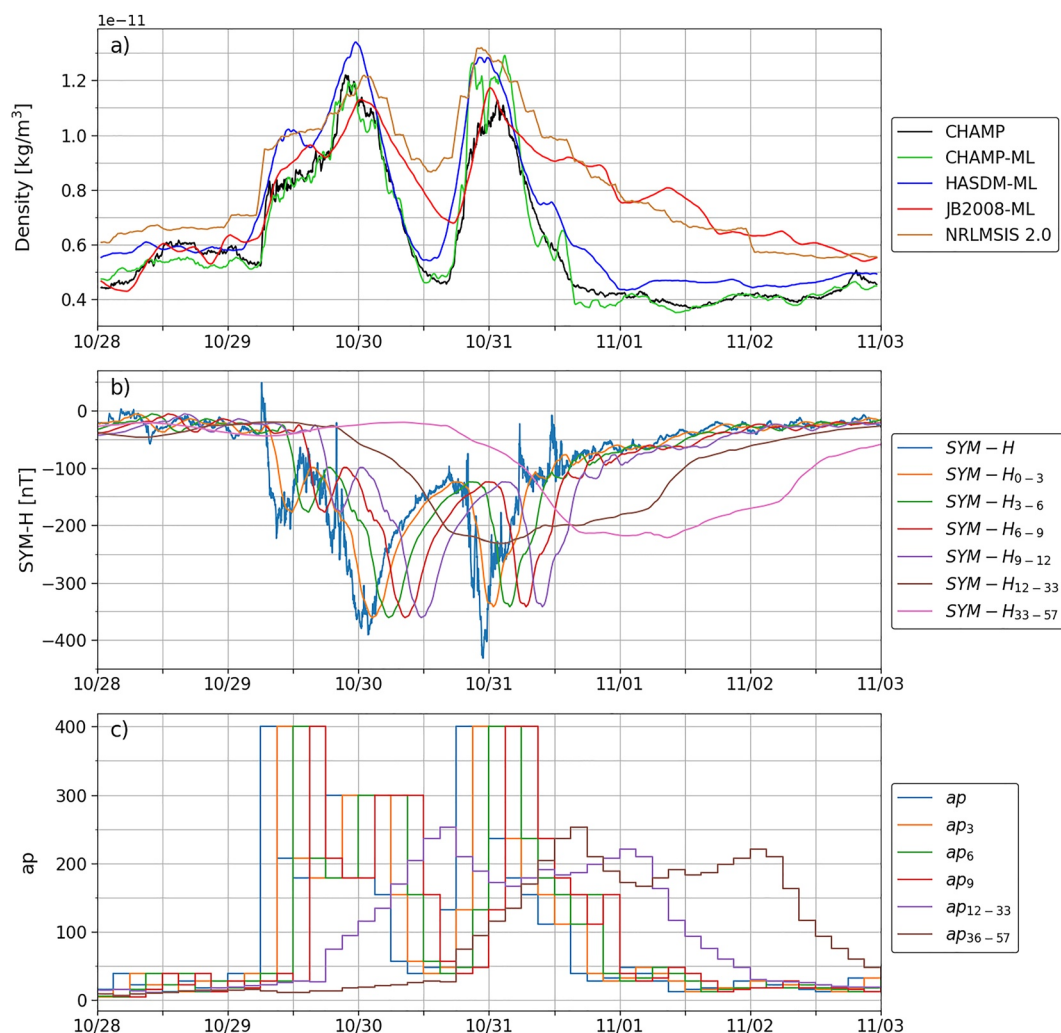
### 2.3. Storm Example

To motivate the work, we evaluated NRLMSIS 2.0 and the ML models from the JB2008, HASDM, and CHAMP data sets during the 2003 Halloween storms. The latter models will be referred to as JB2008-ML, HASDM-ML, and CHAMP-ML, respectively. We used the time series *ap* flag when running NRLMSIS 2.0 in this work. The four models were provided with the true drivers for the 6-day period from 28 October to 3 November 2003 and were compared to the Mehta et al. (2017) CHAMP estimates. For NRLMSIS 2.0 and CHAMP-ML, the predictions were made with the same time cadence and at the specific locations of the satellite, negating the need for further processing. For JB2008-ML and HASDM-ML, we made predictions at the 3-hr intervals used by the original models. Those 3D density grids were then interpolated in space and time in log-scale to the locations of CHAMP. The final step is to take a running average of the along-orbit densities over the 92.5 min orbital period to obtain orbit-averaged densities. This allows us to visualize the general density along the orbit during the storm period (Figure 1).

### 2.4. Time Lag Study

As discussed in Section 1, cooling mechanisms often cause post-storm densities to be anomalously low. For this storm in particular, Lei et al. (2012) noted nearly a 25% decrease in post-storm densities relative to pre-storm levels. In an effort to quantify this mechanism within the original models/data sets, we varied the time histories for *ap* or *SYM-H* independently within the models at four locations listed in Table 2. Table 2 also contains the geomagnetic indices held constant in each model while either *ap* or *SYM-H* were changed. All cases were at a constant solar activity with drivers set to 120. The time inputs were at 0 hr UT and represent the fall equinox (doy = 264), so there are no effects from Earth's tilt.

The models were run with the aforementioned inputs held constant and each of the time series geomagnetic indices increased in magnitude independently. The three models using *ap* step through the 28 discrete values



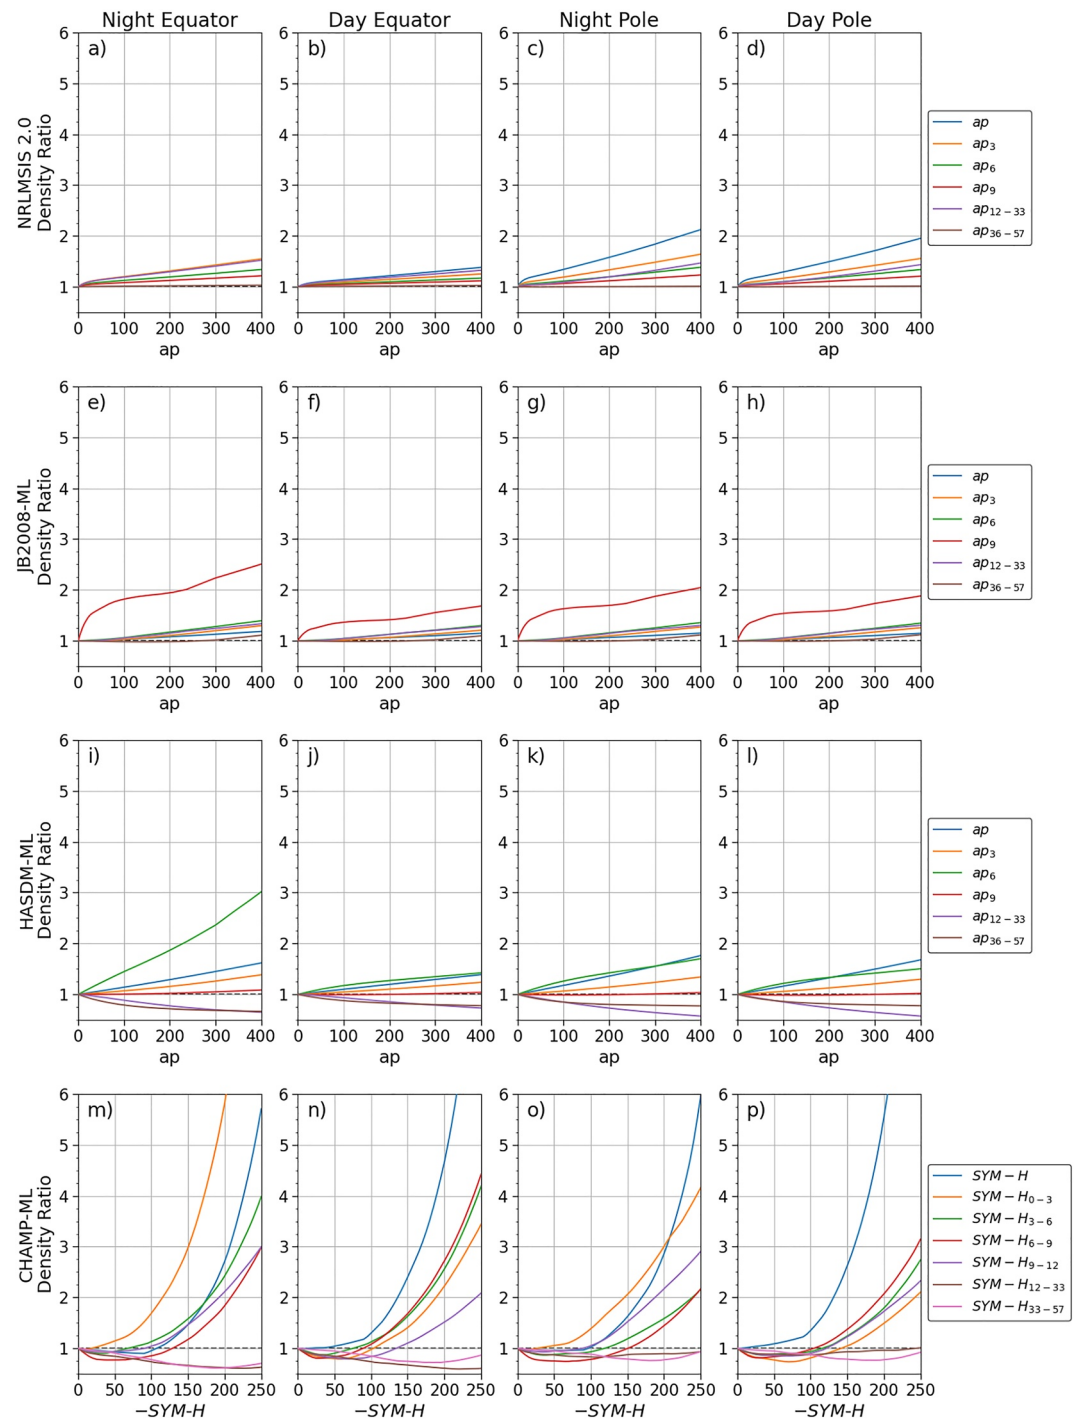
**Figure 1.** Orbit-average density for Naval Research Laboratory Mass Spectrometer and Incoherent Scatter radar, Jacchia-Bowman 2008 Empirical Thermospheric Density Model-Machine Learning (ML), High Accuracy Satellite Drag Model-ML, CHallenging Minisatellite Payload (CHAMP)-ML, and CHAMP (a), the associated  $SYM-H$  (b), and  $ap$  (c) time series inputs.

between 0 and 400. For CHAMP-ML,  $SYM-H$  decreased from 0 to  $-250$  in increments of 5. We then computed the density ratios for each time series index with respect to the density when the index is set to 0. We plot all the curves generated for each location/model (Figure 2).

**Table 2**  
Information for the Time Lag Study

Locations			
Night Equator	Day Equator	Night Pole	Day Pole
LST = 2 hr, LAT = $0^\circ$	LST = 14 hr, LAT = $0^\circ$	LST = 2 hr, LAT = $80^\circ$	LST = 14 hr, LAT = $80^\circ$
Constant Inputs			
NRLMSIS 2.0	JB2008-ML	HASDM-ML	CHAMP-ML
$ap = 56$	$ap = 56$ , $Dst = -50$	$ap = 56$ , $Dst = -50$	$SYM-H = -50$ , $S = 200$
Note. For clarification, LAT is latitude and $S$ refers to both $S_N$ and $S_S$ .			





**Figure 2.** Density ratios for the four models and four locations, as described in Sections 2.2 and 2.4, respectively. The ratios are computed with respect to the particular driver being set to zero.

### 3. Results and Discussion

We first show the error statistics for the three ML models developed in Table 3. These were computed with respect to the original data sets (e.g., JB2008-ML is the error with respect to JB2008 density). Training data are used to fit the model, validation data are used to determine the best model, and the independent test set is used to determine performance; details are in Appendix A.

**Table 3**  
Mean Absolute Error on the Training, Validation, and Test Sets: 100% abs (True-Predicted)/True

Model	Training	Validation	Test
JB2008-ML	5.28%	6.03%	6.63%
HASDM-ML	9.13%	10.46%	10.39%
CHAMP-ML	10.97%	11.60%	11.57%

Table 3 shows that JB2008-ML undoubtedly has the lowest errors, but it is worth noting that it is also the most generalized data set of the three. The SET HASDM density database contains evidence of more complicated processes and its PCA coefficients are more difficult to model as a result (Licata et al., 2021). The CHAMP-ML errors are 1%–2% higher than those of HASDM-ML. CHAMP-ML is the only ML model in this work with location as a predictor. To both visualize the model performance in an operational setting and motivate the remainder of the work, we show the orbit-averaged densities for the three ML models and NRLMSIS 2.0 compared to the Mehta et al. (2017) CHAMP densities for the 2003 Halloween storms in Figure 1.

Figure 1a shows that all models match the timing observed by CHAMP during both storms (29–30 October and 30–31 October). NRLMSIS 2.0 has a tendency to overpredict density throughout this 6-day period, most notably between the two storms and in the recovery phase (02–03 November). This will be explored further with Figure 2. JB2008-ML exhibits similar behavior although it is closer to matching the contraction of the atmosphere between the storms. While both of these models use time histories of  $ap$ , they do not portray any evidence of post-storm cooling. In contrast, HASDM-ML and CHAMP-ML both show significant decreases in density both between and after the storms. Oliveira and Zesta (2019) performed a superposed epoch analysis showing density ratios for CHAMP/GRACE and JB2008 for extreme events which showed that JB2008 is not modeling the satellite observed cooling taking place shortly after these events. In Oliveira et al. (2021), this was expanded to include HASDM. While there were differences between the satellites and HASDM, the sudden cooling was consistent with observations.

Figures 1b and 1c show  $SYM-H$  and  $ap$  time histories, respectively. The increased temporal resolution for  $SYM-H$  is very evident, and the first four averages can inform CHAMP-ML of the recent magnetic disturbances. The last two time history inputs ( $SYM-H_{12-33}$  and  $SYM-H_{33-57}$ ) represent longer-term information with less variation. Immediately following the second storm (around 0600 UTC on 31 October), the last two time history inputs have large magnitudes while the more recent inputs no longer signify a storm. At the same time, the density predicted by CHAMP-ML and observed by the satellite drop abruptly. This behavior is consistent with the observation from Zesta and Oliveira (2019) where CHAMP density experienced sudden decreases following extreme events. The  $ap$  time history is valuable to the other three models, following similar trends to panel (b) but are much more coarse. The results from the time lag study (described in Section 2.4) are displayed in Figure 2.

Figure 2 is informative into what historical information is most important to represent the original data source—JB2008 output, SET HASDM density database, and CHAMP density estimates. NRLMSIS 2.0 is used here as a baseline due to its wide use in the field and use of historical geomagnetic information. There is a fairly linear relationship between the different  $ap$  values and density for NRLMSIS 2.0. In most cases, it considers the most recent  $ap$  to be most important and the least recent  $ap$  to be the least important. There is almost a perfect decay of slopes as it considers information from further in the past. At no point does the density ratio at the four locations drop below 1.00, which represents lower density than the baseline, or  $ap_x = 0$  where  $x$  represents a given time lag or lack thereof. This behavior explains why NRLMSIS 2.0 does not capture the behavior observed by CHAMP in Figure 1. The only drivers for the model that could capture post-storm behavior are the time history  $ap$  inputs. If NRLMSIS 2.0 could model post-storm overcooling, Figures 2a–2d would have to have one or more curves that fall below 1.0 indicating its internal parameters account for this phenomena.

For JB2008-ML, there is virtually no evidence of post-storm cooling being present in the data set. With the exception of the  $ap_9$  curves, there is a fairly linear dependence between  $ap$  and density. Interestingly, JB2008-ML indicates that the strongest relationship between  $ap$  and density has a 9-hr delay. The  $ap_9$  curves are nonlinear for  $ap < 100$  and quite linear for  $ap > 100$ . While there are values for JB2008-ML in Figure 2 that are less than 1.00, they are at most showing a 2% decrease and only at the equatorial locations.

HASDM-ML has a near-linear relationship with  $ap$ , but there is considerable evidence of post-storm cooling seen in Figure 2. At each of the four locations,  $ap_{12-33}$  and  $ap_{36-57}$  have an inverse relationship with density. At the two high-latitude locations,  $ap_{12-33}$  causes the lowest density ratios while  $ap_{36-57}$  causes the lowest density ratios at

the equator. This may be a result of the time delay of the density response at low latitudes relative to the auroral region. In contrast to JB2008-ML panels (e)–(h), HASDM-ML has a strong positive relationship between  $ap_6$  and density with little impact from  $ap_9$ . At the highest levels of activity ( $ap > 300$ ), the current  $ap$  value drives the strongest increase in density at the poles.

CHAMP-ML displays a highly nonlinear relationship between  $SYM-H$  and density. At each location, the relative importance of each input can change significantly; the maximum density ratio for  $SYM-H_{0-3}$  is 10.25 at the nightside equator while it is only 2.10 at the dayside pole. There is strong evidence of post-storm cooling in the CHAMP data set, highlighted by the array of historical  $SYM-H$  drivers causing density ratios below 1.00. The least recent  $SYM-H$  averages have their strongest inverse relationship with density at the equatorial locations while other historical indices demonstrate low density ratios at the polar location. The CHAMP-ML density ratios drop as low as 0.59 and rise as high as 12.34 indicating a more complex relationship between geomagnetic activity and density compared to the other three models in this analysis.

#### 4. Summary

In this work, we demonstrated the use of machine-learned models to quantify the behavior of thermospheric post-storm cooling. We trained ML models on three data sets, JB2008 outputs, the SET HASDM density database, and the Mehta et al. (2017) CHAMP density estimates, to conduct this assessment and used NRLMSIS 2.0 for comparison. All models developed were provided a recent time history (up to 57 hr) of geomagnetic drivers to see if the data suggest that there is evidence of post-storm cooling; the models would need to see whether previous geomagnetic drivers indicate a storm of a given strength that has recently occurred. Using the 2003 Halloween storms as an example (Figure 1), we showed that both NRLMSIS 2.0 and JB2008-ML do not match the sudden cooling seen between and after the two storms by the CHAMP accelerometer. Meanwhile, HASDM-ML and CHAMP-ML both model the general density trends of this storm and display attributes of an abruptly cooled thermosphere.

When considering a historical event, other factors play a role in how the thermosphere behaves. Therefore, we isolated the internal model formulation only as it pertains to recent magnetic perturbations. We all model drivers constant and only vary a single geomagnetic driver at a time:  $ap$  for NRLMSIS 2.0, JB2008-ML, HASDM-ML, and  $SYM-H$  for CHAMP-ML. This showed that NRLMSIS 2.0 and JB2008-ML both did not exhibit any cooling effects as the historical  $ap$  values were raised, which would indicate a strong storm had recently taken place. In fact, the most important historical driver to JB2008-ML was the 9-hr prior  $ap$  which resulted in density ratios nearly twice that of any other driver. This is due to a 9-hr offset in the calculation of the JB2008 parameter DTC based upon  $ap$  use (Bowman and Tobiska, private communication, 2020).

HASDM-ML was most strongly driven by the current and 6-hr prior  $ap$  for thermospheric expansion while increases in  $ap_{12-33}$  and  $ap_{36-57}$  resulted in densities as low as 57% of the baseline magnitude. CHAMP-ML was the only model to indicate a nonlinear relationship between density and geomagnetic activity. Depending on the location,  $SYM-H$  or  $SYM-H_{0-3}$  drove the largest density ratios, significantly more than any other model. In terms of cooling, CHAMP-ML showed that at  $SYM-H > -100$  nT, many of the recent drivers caused a density ratio less than 1.00. As the index was made more negative, the least recent drivers caused the lowest density ratios, particularly at low latitudes.

Time lagged components of geomagnetic indices used in standard space weather models were evaluated with ML regression models to elucidate their significance in representing the observed behavior of thermospheric cooling phenomena in observed data sets. The “hidden” internal structure of the ML regression models are difficult to examine, but an analysis like this can inform us what the overall formulation the model has with respect to its drivers. This method can be used to gain a better understanding of the model and to investigate complex space weather data sets for improved future empirical and physics-based modeling.



## 5. Recommendations

This study used existing models from Licata et al. (2022) and Licata and Mehta (2022) which were developed with the goal of forecasting in mind. This led to the decision to include the geomagnetic index time series as ML model inputs. Direct measurements of NO or CO<sub>2</sub> emissions were not considered, as at present such data are not available with sufficient lead time for forecasting or nowcasting. Nevertheless, a future worthwhile study could include testing post-storm prediction capabilities of these models with and without these additional drivers. If a future mission could provide real-time measurements of NO emissions at multiple locations, then these could be used in a forecasting model similar to HASDM-ML or CHAMP-ML.

## Appendix A: Data Splitting

When developing a ML model, one must carefully consider how to split the data into training, validation, and test sets. These models can have thousands or even millions of parameters, so it is critical to ensure the analysis is valid and that the model has not just memorized the training data. The goal is to provide enough diverse data in the training set for the model to learn the proper relationship between the inputs and outputs. The validation and test sets are chosen to also contain diverse conditions. The validation set is used to determine the model performance—during and after training—on data not used for fitting. However, to choose the best version of the model based on its training and validation performance, there must be another independent data set for model evaluation; this is called the test set.

For JB2008-ML and HASDM-ML, the data splitting is straightforward. Looking at the solar and geomagnetic drivers, we split the data by years trying to place various portions of the solar cycle in each set. The breakdown is displayed in Table A1. This results in 12 years being used for training (60%), 4 years being used for validation (20%), and 4 years being used for test (20%).

For CHAMP-ML, the combination of having location as an input and not having a full solar cycle of data makes the data splitting more complex. Using year-long periods (as in Table A1) leaves out an entire range of possible altitudes and solar drivers from training. We work around this by repeating the following scheme: Eight consecutive weeks for training, 1 week for validation, and 1 week for test. Even though there are only two weeks between training segments, the 10-s cadence of the measurements creates a gap of over 120,000 samples, ensuring the model is not just simply interpolating for the validation and test sets. While only 60% of data are in the JB2008-ML and HASDM-ML training sets, the spatiotemporally limited CHAMP data set requires more training data to prevent overfitting.

**Table A1**  
*Training, Validation, and Test Split for JB2008-Machine Learning (ML) and HASDM-ML by Year*

Set	Training	Validation	Test
<b>Years</b>	2000, 2001, 2003, 2004, 2007, 2009, 2011, 2012, 2013, 2015, 2016, 2019	2005, 2008, 2014, 2017	2002, 2006, 2010, 2018

## Appendix B: Additional Model Development Details

### B1. Keras Tuner

As discussed in Section 2.2, we use Keras Tuner to identify an architecture for each data set. The hyperparameter space we defined for the tuners (one for each data set) is displayed in Table B1.

For JB2008-ML and HASDM-ML, the tuner is provided the entire training and validation sets due to the relatively small size of the data sets. Therefore, once the tuner returns the best models, each trained for 100 training iterations or epochs, there is no need for further training. In fact, training after tuning did not yield improved models in these two cases. The CHAMP-ML training and validation sets have over 20 million and 2 million samples, respectively. Therefore, we only supply the tuner with 1 million random samples from the training set and 200,000 random samples from the validation set. This provides the tuner with ample information to determine an adequate architecture from which we can develop the full model. This full model will use the full training and validation sets.

**Table B1**  
*Hyperparameter Search Space for the Three Machine Learning Models*

Parameter	Choices
Number of Hidden Layers	1–10
Neurons	min = 64, max = 1,024, step = 4
Activations	relu, softplus, tanh, sigmoid, softsign, selu, elu, linear
Dropout	min = 0.10, max = 0.60, step = 0.01
Optimizer	RMSprop, Adam, Nadam, Adadelta, Adagrad

## B2. Loss Function and Custom Layer

The loss function plays an important role in the tuning and training process. It is the metric that the fitting function tries to minimize or maximize through the weight updates. For this study, a mean square error (MSE) loss function would suffice since the goal is to develop a model with minimal error with respect to the respective data set. However, we used models from previous work, Licata et al. (2022), which used the negative logarithm of predictive density (NLPD) loss function:

$$NLPD(y, \mu, \sigma) = \frac{(y - \mu)^2}{2\sigma^2} + \frac{\ln(\sigma^2)}{2} + \frac{\ln(2\pi)}{2} \quad (B1)$$

where  $y$ ,  $\mu$ , and  $\sigma$  are the ground truth, mean prediction, and standard deviation of the prediction, respectively. In practice, the factor of  $1/2$  and the last term can be removed. There are no ground truth values for  $\sigma$  for these data sets, but we are able to directly predict it since it is a standalone term in Equation B1. This is accomplished by using twice the output neurons as there are outputs where they represent the mean and standard deviation of the model uncertainty. The output neurons representing the mean use a linear activation function, while those representing the standard deviation use the "softplus" activation function. The softplus activation and its derivative, the sigmoid function, are defined as following:

$$f(x) = \ln(1 + e^x) \quad f'(x) = \frac{e^x}{1 + e^x} \quad (B2)$$

with  $x$  being the weighted sum of the neurons inputs. Licata et al. (2022) found that for the SET HASDM density database, the use of NLPD as opposed to MSE did not result in significantly different prediction errors.

## B3. Batch Size for Tuners and Further Training

Batch size is the number of samples used to average the loss during fitting. With too small of a batch size, the model can update its weights often to specific batches of the data making it unstable. With too large of a batch size, the model can have trouble learning due to the over-generalization of the losses. Choosing a batch size for a given application requires trial and error as there is no universal choice. Table B2 shows the batch sizes used for the tuners and subsequent training for each data set.

Batch sizes in Table B2 were chosen from trial and error. The JB2008-ML and HASDM-ML batch size is the smallest partially due to the size of the data set (35,064 training samples). For the CHAMP-ML tuner, we use 1 million samples and therefore have a larger batch size. The batch size for training the final CHAMP-ML model using the best architecture from the tuner is significantly larger, because we have over 20 million training samples and found this leads to the most stable training process.

**Table B2**  
*Batch Sizes Used to Tune and Train the Three Machine Learning Models in This Work*

	JB2008-ML/HASDM-ML	CHAMP-ML
Tuner	$2^8 = 256$	$2^{12} = 4,096$
Further Training	N/A	$2^{17} = 131,072$

## Data Availability Statement

Requests can be submitted for full access to the SET HASDM density database at <https://spacewx.com/hasdm/> and all reasonable requests for scientific research are accepted as explained in the rules of road document on the website. The historical space weather indices used in this study can be found at the following links:  $F_{10.7}$ : <https://www.spaceweather.gc.ca/forecast-previous/solar-solaire/solarflux/sx-en.php>,  $ap$ : <https://doi.org/10.5880/Kp.0001>,  $Dst$ : <http://wdc.kugi.kyoto-u.ac.jp/dst/dir/>, and  $SYM-H$ : <http://wdc.kugi.kyoto-u.ac.jp/aeasy/index.html>. The remaining solar indices and proxies can be found at <https://spacewx.com/jb2008/> in the SOLFSMY.TXT file. The Weimer (2019) Pointing flux data can be accessed at <https://doi.org/10.5281/zenodo.3525166>. Free, onetime only registration is required to access the historical data while nowcasts and forecasts are provided by SET as a data service from [data@spacewx.com](mailto:data@spacewx.com). CHALLENGING Minisatellite Payload and GRACE position data were obtained from the measurements presented by Mehta et al. (2017) at <https://tinyurl.com/densitysets>. Programs and files used for ML model development and analysis can be found at <https://doi.org/10.5281/zenodo.7075028>.

## Acknowledgments

PMM gratefully acknowledges support under NSF CAREER award no. 2140204 and NSF ANSWERS award no. 2149747 (subaward to WVU from Rutgers). All authors acknowledge support by NASA Grant No. 80NSSC20K1362 to Virginia Tech under the Space Weather Operations 2 Research Program, with subcontracts to WVU and SET. Douglas Drob was supported by NASA interagency agreement 80HQTR20T0081 with the Naval Research Laboratory. We would like to thank Space Weather Canada for providing and maintaining solar radio emission data, GFZ Potsdam for supplying  $ap$  archives, and the World Data Center for Geomagnetism in Kyoto for providing  $Dst$  and  $SYM-H$  data. The authors would like to acknowledge DLR for their work on the CHAMP mission along with GFZ Potsdam for managing the data.

## References

- Bettadpur, S. (2012). *Gravity recovery and climate experiment: Product specification document*. GRACE 327-720, CSR-GR-03-02. Centre for Space Research, The University of Texas. Retrieved from <https://podaac.jpl.nasa.gov/GRACE>
- Bowman, B., Tobiska, W. K., Marcos, F., Huang, C., Lin, C., & Burke, W. (2008). A new empirical thermospheric density model JB2008 using new solar and geomagnetic indices. AIAA 2008-6438. In *AIAA/AAS astrodynamics specialist conference*. <https://doi.org/10.2514/6.2008-6438>
- Bruinsma, S., & Biancale, R. (2003). Total densities derived from accelerometer data. *Journal of Spacecraft and Rockets*, 40(2), 230–236. <https://doi.org/10.2514/2.3937>
- Chellapilla, K., Puri, S., & Simard, P. (2006). High performance convolutional neural networks for document processing. In G. Lorette & U. deRennes (Eds.), *Tenth international workshop on frontiers in handwriting recognition* (Vol. 1). Suvisoft.
- Covington, A. E. (1948). Solar noise observations on 10.7 centimeters. *Proceedings of the IRE*, 36(4), 454–457. <https://doi.org/10.1109/JRPROC.1948.234598>
- Doornbos, E. (2012). *Producing density and crosswind data from satellite dynamics observations* (pp. 91–126). Springer Berlin Heidelberg. [https://doi.org/10.1007/978-3-642-25129-0\\_4](https://doi.org/10.1007/978-3-642-25129-0_4)
- Emmert, J. (2015). Thermospheric mass density: A review. *Advances in Space Research*, 56(5), 773–824. <https://doi.org/10.1016/j.asr.2015.05.038>
- Emmert, J. T., Drob, D. P., Picone, J. M., Siskind, D. E., Jones, M., Jr., Mlynarczyk, M. G., et al. (2021). NRLMSIS 2.0: A whole-atmosphere empirical model of temperature and neutral species densities. *Earth and Space Science*, 8(3), e2020EA001321. <https://doi.org/10.1029/2020EA001321>
- Florios, K., Kontogiannis, I., Park, S.-H., Guerra, J. A., Benvenuto, F., Bloomfield, D. S., & Georgoulis, M. K. (2018). Forecasting solar flares using magnetogram-based predictors and machine learning. *Solar Physics*, 293(2), 1–42. <https://doi.org/10.1007/s11207-018-1250-4>
- Fuller-Rowell, T. J., Codrescu, M. V., Roble, R. G., & Richmond, A. D. (1997). *How does the thermosphere and ionosphere react to a geomagnetic storm?* (pp. 203–225). American Geophysical Union (AGU). <https://doi.org/10.1029/GM098p0203>
- Hedin, A. E. (1987). MSIS-86 thermospheric model. *Journal of Geophysical Research*, 92(A5), 4649–4662. <https://doi.org/10.1029/JA092iA05p04649>
- Iyemori, T. (1990). Storm-time magnetospheric currents inferred from mid-latitude geomagnetic field variations. *Journal of Geomagnetism and Geoelectricity*, 42(11), 1249–1265. <https://doi.org/10.5636/jgg.42.1249>
- Jiao, Y., Hall, J. J., & Morton, Y. T. (2017). Automatic equatorial GPS amplitude scintillation detection using a machine learning algorithm. *IEEE Transactions on Aerospace and Electronic Systems*, 53(1), 405–418. <https://doi.org/10.1109/TAES.2017.2650758>
- Knipp, D. J., Pette, D. V., Kilcommons, L. M., Isaacs, T. L., Cruz, A. A., Mlynarczyk, M. G., et al. (2017). Thermospheric nitric oxide response to shock-led storms. *Space Weather*, 15(2), 325–342. <https://doi.org/10.1002/2016SW001567>
- Kockarts, G. (1980). Nitric oxide cooling in the terrestrial thermosphere. *Geophysical Research Letters*, 7(2), 137–140. <https://doi.org/10.1029/GL007i002p00137>
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems* (Vol. 25). Curran Associates, Inc.
- Lei, J., Burns, A. G., Thayer, J. P., Wang, W., Mlynarczyk, M. G., Hunt, L. A., et al. (2012). Overcooling in the upper thermosphere during the recovery phase of the 2003 October storms. *Journal of Geophysical Research*, 117(A3). <https://doi.org/10.1029/2011JA016994>
- Lei, J., Wang, W., Burns, A. G., Zhang, S.-R., & Dang, T. (2021). Comments on “poststorm thermospheric NO overcooling?” By Mikhailov and Perrone (2020). *Journal of Geophysical Research: Space Physics*, 126(4), e2020JA027992. <https://doi.org/10.1029/2020JA027992>
- Licata, R. J., & Mehta, P. M. (2022). Uncertainty quantification techniques for space weather modeling: Thermospheric density application.
- Licata, R. J., Mehta, P. M., Tobiska, W. K., Bowman, B. R., & Pilinski, M. D. (2021). Qualitative and quantitative assessment of the SET HASDM database. *Space Weather*, 19(8), e2021SW002798. <https://doi.org/10.1029/2021SW002798>
- Licata, R. J., Mehta, P. M., Tobiska, W. K., & Huzurbazar, S. (2022). Machine-learned HASDM thermospheric mass density model with uncertainty quantification. *Space Weather*, 20(4), e2021SW002915. <https://doi.org/10.1029/2021SW002915>
- Liu, H., Lühr, H., Henize, V., & Köhler, W. (2005). Global distribution of the thermospheric total mass density derived from CHAMP. *Journal of Geophysical Research*, 110(A4). <https://doi.org/10.1029/2004JA010741>
- March, G., Doornbos, E., & Visser, P. (2019). High-fidelity geometry models for improving the consistency of CHAMP, GRACE, GOCE and Swarm thermospheric density data sets. *Advances in Space Research*, 63(1), 213–238. <https://doi.org/10.1016/j.asr.2018.07.009>
- Mehta, P. M., & Linares, R. (2017). A methodology for reduced order modeling and calibration of the upper atmosphere. *Space Weather*, 15(10), 1270–1287. <https://doi.org/10.1002/2017SW001642>
- Mehta, P. M., Linares, R., & Sutton, E. K. (2018). A quasi-physical dynamic reduced order model for thermospheric mass density via Hermitian space-dynamic mode decomposition. *Space Weather*, 16(5), 569–588. <https://doi.org/10.1029/2018SW001840>
- Mehta, P. M., Walker, A. C., Sutton, E. K., & Godinez, H. C. (2017). New density estimates derived using accelerometers on board the CHAMP and GRACE satellites. *Space Weather*, 15(4), 558–576. <https://doi.org/10.1002/2016SW001562>

- Mikhailov, A. V., & Perrone, L. (2020). Poststorm thermospheric NO overcooling? *Journal of Geophysical Research: Space Physics*, 125(1), e2019JA027122. <https://doi.org/10.1029/2019JA027122>
- Mlynczak, M., Martin-Torres, F. J., Russell, J., Beaumont, K., Jacobson, S., Kozyra, J., et al. (2003). The natural thermostat of nitric oxide emission at 5.3  $\mu\text{m}$  in the thermosphere observed during the solar storms of April 2002. *Geophysical Research Letters*, 30(21), 2100. <https://doi.org/10.1029/2003GL017693>
- O'Malley, T., Bursztein, E., Long, J., Chollet, F., Jin, H., Invernizzi, L., et al. (2019). Keras tuner. Retrieved from <https://github.com/keras-team/keras-tuner>
- Oliveira, D. M., & Zesta, E. (2019). Satellite orbital drag during magnetic storms. *Space Weather*, 17(11), 1510–1533. <https://doi.org/10.1029/2019SW002287>
- Oliveira, D. M., Zesta, E., Mehta, P. M., Licata, R. J., Pilinski, M. D., Tobiska, W. K., & Hayakawa, H. (2021). The current state and future directions of modeling thermosphere density enhancements during extreme magnetic storms. *Frontiers in Astronomy and Space Sciences*, 8. <https://doi.org/10.3389/fspas.2021.764144>
- Reigber, C., Lühr, H., & Schwintzer, P. (2002). Champ mission status. *Advances in Space Research*, 30(2), 129–134. [https://doi.org/10.1016/S0273-1177\(02\)00276-4](https://doi.org/10.1016/S0273-1177(02)00276-4)
- Russell, J. M. I., Mlynczak, M. G., Gordley, L. L., Tansock, J., & Esplin, R. (1999). An overview of the SABER experiment and preliminary calibration results. *Space Dynamics Lab Publications*, 114. <https://doi.org/10.1117/12.366382>
- Siciliano, F., Consolini, G., Tozzi, R., Gentili, M., Giannattasio, F., & De Michelis, P. (2021). Forecasting SYM-H index: A comparison between long short-term memory and convolutional neural networks. *Space Weather*, 19(2), e2020SW002589. <https://doi.org/10.1029/2020SW002589>
- Storz, M. F., Bowman, B. R., Branson, M. J. I., Casali, S. J., & Tobiska, W. K. (2005). High accuracy satellite drag model (HASDM). *Advances in Space Research*, 36(12), 2497–2505. <https://doi.org/10.1016/j.asr.2004.02.020>
- Sutton, E. K. (2008). *Effects of solar disturbances on the thermosphere densities and winds from CHAMP and GRACE satellite accelerometer data*. Ph.D. thesis, University of Colorado at Boulder. Retrieved from <https://ui.adsabs.harvard.edu/abs/2008PhDT...87S>
- Tobiska, W. K., Bouwer, S. D., & Bowman, B. R. (2008). The development of new solar indices for use in thermospheric density modeling. *Journal of Atmospheric and Solar-Terrestrial Physics*, 70(5), 803–819. <https://doi.org/10.1016/j.jastp.2007.11.001>
- Tobiska, W. K., Bowman, B. R., Bouwer, D., Cruz, A., Wahl, K., Pilinski, M., et al. (2021). *The SET HASDM density database*. *Space Weather*. <https://doi.org/10.1029/2020SW002682>. e2020SW002682
- Weimer, D. (2019). Improving neutral density predictions using exospheric temperatures calculated on a geodesic. Polyhedral Grid [Data set], <https://doi.org/10.5281/zenodo.3525166>
- Weimer, D. R. (2005a). Improved ionospheric electrodynamic models and application to calculating Joule heating rates. *Journal of Geophysical Research*, 110(A5), A05306. <https://doi.org/10.1029/2004JA010884>
- Weimer, D. R. (2005b). Predicting surface geomagnetic variations using ionospheric electrodynamic models. *Journal of Geophysical Research*, 110(A12), A12307. <https://doi.org/10.1029/2005JA011270>
- Wing, S., Johnson, J. R., Jen, J., Meng, C.-I., Sibeck, D. G., Bechtold, K., et al. (2005). Kp forecast models. *Journal of Geophysical Research*, 110(A4). <https://doi.org/10.1029/2004JA010500>
- Zesta, E., & Oliveira, D. M. (2019). Thermospheric heating and cooling times during geomagnetic storms, including extreme events. *Geophysical Research Letters*, 46(22), 12739–12746. <https://doi.org/10.1029/2019GL085120>