



OPEN

Uncertainty quantification techniques for data-driven space weather modeling: thermospheric density application

Richard J. Licata^{1,2}✉ & Piyush M. Mehta^{1,2}

Machine learning (ML) has been applied to space weather problems with increasing frequency in recent years, driven by an influx of in-situ measurements and a desire to improve modeling and forecasting capabilities throughout the field. Space weather originates from solar perturbations and is comprised of the resulting complex variations they cause within the numerous systems between the Sun and Earth. These systems are often tightly coupled and not well understood. This creates a need for skillful models with knowledge about the confidence of their predictions. One example of such a dynamical system highly impacted by space weather is the thermosphere, the neutral region of Earth's upper atmosphere. Our inability to forecast it has severe repercussions in the context of satellite drag and computation of probability of collision between two space objects in low Earth orbit (LEO) for decision making in space operations. Even with (assumed) perfect forecast of model drivers, our incomplete knowledge of the system results in often inaccurate thermospheric neutral mass density predictions. Continuing efforts are being made to improve model accuracy, but density models rarely provide estimates of confidence in predictions. In this work, we propose two techniques to develop nonlinear ML regression models to predict thermospheric density while providing robust and reliable uncertainty estimates: Monte Carlo (MC) dropout and direct prediction of the probability distribution, both using the negative logarithm of predictive density (NLPD) loss function. We show the performance capabilities for models trained on both local and global datasets. We show that the NLPD loss provides similar results for both techniques but the direct probability distribution prediction method has a much lower computational cost. For the global model regressed on the Space Environment Technologies High Accuracy Satellite Drag Model (HASDM) density database, we achieve errors of approximately 11% on independent test data with well-calibrated uncertainty estimates. Using an in-situ CHALLENGING Minisatellite Payload (CHAMP) density dataset, models developed using both techniques provide test error on the order of 13%. The CHAMP models—on validation and test data—are within 2% of perfect calibration for the twenty prediction intervals tested. We show that this model can also be used to obtain global density predictions with uncertainties at a given epoch.

Low Earth orbit (LEO) will see the addition of tens of thousands of satellites in the coming years as private companies are developing mega constellations for the new space economy¹. This congestion of certain orbital regimes increases the likelihood of a future collision between two objects. Satellite collisions can create debris clouds consisting of thousands of objects large enough to pose significant threats to other space assets. The 2009 Iridium-Cosmos collision resulted in approximately 2300 observable debris objects, 65% of which remained in orbit 7 years later². Debris objects created by collisions or weapons tests can catapult into highly elliptical orbits which pose a danger to satellites in multiple orbital regimes³.

In an effort to prevent these events from occurring, objects are continuously tracked, and their trajectories are predicted. However, uncertainties play a large role in the prediction of future satellite positions. In LEO, atmospheric drag is the largest single source of uncertainty mainly due to an incomplete understanding of the thermosphere. Variations in the thermosphere are connected to temperature changes, as the atmosphere expands and contracts. Solar extreme ultraviolet (EUV) and far ultraviolet (FUV) irradiance are the primary heating

¹Department of Mechanical and Aerospace Engineering, West Virginia University, Morgantown, WV 26506, USA.

²These authors contributed equally: Richard J. Licata and Piyush M. Mehta. ✉email: rjlicata@mix.wvu.edu

sources⁴. This absorption of solar irradiance provides the baseline thermospheric mass density⁵. The effects of solar emissions are well-represented by various solar indices and proxies⁶.

Solar irradiance is generally a long-term variation while the solar wind drives more rapid changes in the thermosphere. Mass and energy from the sun—manifested as the solar wind—travel through space and interact with the near-Earth geospace environment. Certain events (e.g. coronal mass ejections) send massive amounts of energy and mass that result in significant increases in thermospheric density. Energy, and therefore density, enhancements first appear in the auroral zone (high latitudes) and propagate towards the equator in the form of traveling atmospheric disturbances⁷. Geomagnetic storms are a particularly difficult phenomena to model and our current density models carry high uncertainty during these periods^{8,9}.

Satellite accelerometers have provided a unique insight into the thermosphere with high fidelity in-situ measurements, particularly during storms¹⁰. Accelerations caused by non-drag sources (e.g. gravity and solar radiation pressure) are modeled out allowing the isolation of drag acceleration that is then used to estimate mass density^{11–15}. Drag acceleration is given as

$$\vec{a}_{\text{drag}} = -\frac{1}{2}\rho\frac{c_D A}{m}v_{\text{rel}}^2\frac{v_{\text{rel}}}{|v_{\text{rel}}|} \quad (1)$$

where \vec{a}_{drag} is the drag acceleration, ρ is local mass density, c_D is the satellite drag coefficient, A is the cross-sectional area, m is the satellite mass, and v_{rel} is the relative velocity of the satellite with respect to the rotating atmosphere. With an estimate for drag acceleration, the density can be estimated, assuming adequate knowledge of the drag coefficient and cross-sectional area given the satellite orientation. Density estimates obtained through this method are considered ground truth and often used for model validation.

Accelerometer and orbit-derived densities have been used frequently in developing empirical models^{16–18}. Furthermore, they have been used in data assimilation schemes to make corrections to background models, either through observed orbital drag data¹⁹ or two-line element data²⁰. The most prominent integration of real-time data for neutral density modeling is the High Accuracy Satellite Drag Model's (HASDM) Dynamic Calibration of the Atmosphere (DCA). This uses observed satellite data to make corrections to a background empirical density model²¹.

Even with these improvements, density models have high errors, and we generally use them without any knowledge of their confidence given the conditions. Until recently, no thermospheric density models—whether physics-based or empirical—provided estimates of uncertainty. Bruinsma et al. developed an uncertainty-based version of DTM2020 using polynomials to describe the 1 – sigma uncertainties as a function of the inputs²². Licata et al. used MC dropout to obtain uncertainty estimates for a global density modeling application with good calibration, providing baseline performance²³. In this work, we leverage machine learning (ML) to generate predictive density models for the thermosphere that also provide robust and reliable uncertainty estimates. This is done for both a global and local datasets using two methods: Monte Carlo (MC) dropout and a direct prediction of the probability distribution (referred to primarily as direct probability).

We first outline the data and methods used for model development and analysis. Then, we use artificial data to demonstrate the techniques. We move to the results for modeling with the global density dataset using the two uncertainty techniques and perform a similar analysis for the models developed on local measurements. We also look at the global prediction capabilities of the model developed with in-situ data, and we compare the evaluation times of both uncertainty methods.

SET HASDM density database. The High Accuracy Satellite Drag Model (HASDM) is the operational thermospheric density framework used by the United States Space Force (USSF)²¹. By improving the density correction techniques presented by Marcos et al.²⁴ and Nazerenko et al.²⁵, HASDM modifies 13 global temperature coefficients to make real-time corrections to the Jacchia-Bowman 2008 Empirical Thermospheric Density Model (JB2008)¹⁷. Its Dynamic Calibration of the Atmosphere (DCA) algorithm ingests data from calibration satellites that are distributed in altitude between 190 and 900 km—most are between 300 and 600 km²⁶. As the algorithm provides corrections to JB2008, HASDM provides global measurements on a $24 \times 19 \times 27$ grid. For additional information on HASDM, the reader is referred to Storz et al.²¹.

While HASDM is highly desired due to its real-time data assimilation, it is a proprietary model that is inaccessible to researchers and operators. Space Environment Technologies (SET) is the contractor responsible for validating HASDM outputs on a weekly basis, and they recently released HASDM validation archives from 2000 to 2020 covering close to two full solar cycles providing good statistical coverage. This data release constitutes the SET HASDM density database²⁷. With a 3-h cadence, the database contains 58,440 global HASDM outputs. Each output has a resolution of 15° longitude, 10° latitude, and 25 km altitude ranging from 175 to 825 km. For further details on the SET HASDM density database and the validation process, the reader is referred to Tobiska et al.²⁷.

Satellite accelerometer density estimates. CHALLENGING Minisatellite Payload (CHAMP) was launched in mid-2000 to study Earth's gravity and magnetic fields²⁸. Its orbit is nearly polar with an inclination of 87.3° providing adequate global coverage, and it began at 460 km in altitude. CHAMP was in orbit until 2010 when it stopped providing measurements at an altitude of approximately 300 km. The long mission lifetime covered nearly a solar cycle, providing measurements in solar maximum across many strong geomagnetic storms and through the following solar minimum. Mehta et al. used higher fidelity satellite geometry and improved gas-surface interaction models to scale the CHAMP density estimates of Sutton^{29–32}. The density dataset starts on January 1, 2002 and ends on February 22, 2010 with a 10-s cadence. The CHAMP dataset is prime for demonstration as spatiotemporally limited in-situ datasets are common in the space weather field. This type of model can be built upon with the addition of density estimates from other satellites, displayed in Fig. 1.




Figure 1. Timeline of satellites with onboard accelerometers from 2000 to 2018 with bold lines representing the satellites’ mean altitude. e-POP does not contain an accelerometer. The bottom panel shows the corresponding F_{10} values indicating solar activity. The authors gratefully acknowledge Dr. Eelco Doornbos for providing this figure.

The addition of all satellites shown in Fig. 1 would significantly expand the altitude coverage of the in-situ density dataset. The CHAMP dataset used in this work has a 160 km altitude range and does not span a full solar cycle. Integrating the density datasets of Gravity Recovery and Climate Experiment (GRACE)³³, Gravity Field and Steady-State Ocean Circulation Explorer (GOCE)³⁴, and Swarm³⁵ would provide a thorough altitude coverage of approximately 220–550 km and span from 2001 to present day. The Enhanced Polar Outflow Probe (e-POP)³⁶ is a payload on Cascade, Smallsat and Ionospheric Polar Explorer (CASSIOPE) and its density estimates can be obtained through the processing of its Global Navigation Satellite System (GNSS) receivers^{37,38}. However, there is still much active research related to the proper combination of different satellite density datasets^{39–41}. Therefore, we proceed with the standalone CHAMP dataset for demonstration.

Density model drivers. JB2008 uses four solar indices/proxies as drivers for solar activity. F_{10} —more completely referred to as $F_{10.7}$ —represents 10.7 cm solar radio flux and is a reliable proxy for solar EUV heating. S_{10} is an index for the integrated 26–34 nm solar EUV emission. The M_{10} proxy is a surrogate for FUV photospheric 160-nm Schumann-Runge Continuum emissions. Y_{10} is a hybrid index that measures solar coronal X-ray emissions during solar maximum and Lyman- α emissions during solar minimum. The S_{10} , M_{10} , and Y_{10} indices and proxies are not related to the 10.7 cm wavelength, but they are converted to F_{10} units—solar flux units (sfu)—through linear regression. JB2008 also uses the 81-day centered averages for all four solar drivers. This is indicated by the “81c” subscript. Additional information on these solar drivers is provided by Tobiska et al.⁶.

To model geomagnetic activity, JB2008 uses a combination of a_p and Dst . The a_p index represents global geomagnetic activity with a 3-h cadence. While it is widely used in density models, it is limited by the low-latitude range of the measurements and its discrete range of 28 values. Dst is an index driven by the ring current strength in the inner magnetosphere⁴². When Dst_{min} is below -75 nT, JB2008 shifts to using Dst as it improves storm-time performance¹⁷. The EXEMPLAR (EXospheric TEMperatures on a PoLyhedrAl gRid) model and EXEMPLAR-ML use Poynting flux totals in the northern and southern hemispheres— S_N and S_S , respectively^{43,44}. Poynting flux represents electrodynamic energy flowing into the upper atmosphere. The a_p and Dst indices have 3-h and 1-h cadences, respectively. Therefore, their use in a high-cadence model would not be advised. The geomagnetic index used to replace a_p and Dst in the CHAMP model is $SYM-H$, the longitudinally symmetric component of the magnetic field disturbances^{45,46}. $SYM-H$ is available with a 1-min cadence.

The input sets for the HASDM and CHAMP models are shown in Table 1. The transformed time inputs t_1 – t_4 are defined in Eq. (2), and the transformed local solar time (LST) inputs are defined in Eq. (3). These transformations are performed to make the time and location inputs continuous.

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

$$LST_1 = \sin\left(\frac{2\pi LST}{24}\right) \quad LST_2 = \cos\left(\frac{2\pi LST}{24}\right) \quad (3)$$

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

Table 1. List of inputs for both models.

In Table 1, LST , LAT , and ALT are the local time, latitude, and altitude of the satellite, respectively. The “A” subscript for the geomagnetic indices refers to the daily average. The numerical subscripts for these indices refer to the value that many hours prior to prediction epoch. The combination of numbers refers to the average of that index over that many hours prior to epoch. The HASDM input set originates from Licata et al.²³.

Methodology

Machine learning. *Principal component analysis.* Principal component analysis (PCA), also referred to as Empirical Orthogonal Function (EOF) analysis, has been widely used in thermospheric mass density applications. It has been applied to satellite accelerometer datasets (as described in “[Satellite accelerometer density estimates](#)”) to identify dominant modes of variability in the thermosphere^{15,47,48}. PCA is also used in the field for dimensionality reduction as part of a reduced-order model (ROM)^{49–51}. The HASDM dataset has 12,312 model outputs each epoch which makes uncertainty quantification (UQ) infeasible. Therefore, we apply PCA to the dataset for ROM development with the goal of UQ. PCA is an eigendecomposition technique that maximizes variance, determining uncorrelated linear combinations of data^{52,53}. We first take the common logarithm (\log_{10}) of the density values to reduce the data’s variance then remove the spatial mean. PCA decomposes the data, separating the spatial and temporal variations such that

$$\mathbf{x}(\mathbf{s}, t) = \bar{\mathbf{x}}(\mathbf{s}) + \tilde{\mathbf{x}}(\mathbf{s}, t) \quad \text{and} \quad \tilde{\mathbf{x}}(\mathbf{s}, t) = \sum_{i=1}^r \alpha_i(t) U_i(\mathbf{s}) \quad (4)$$

In Eq. (4), $\mathbf{x}(\mathbf{s}, t)$ is the log density from HASDM, $\bar{\mathbf{x}}$ is the spatial mean, and $\tilde{\mathbf{x}}$ is the variation about the mean. $\alpha_i(t)$ are temporal PCA coefficients and U_i are orthogonal modes—also called basis functions. The choice order of truncation ($r = 10$) was chosen from analyses in previous work as it allows for $\sim 90\%$ of the variance to be captured and only results in $< 3\%$ truncation error^{23,54}. The orthogonal modes are derived through

$$\mathbf{X} = U \Sigma V^T \quad \text{where} \quad \mathbf{X} = \begin{bmatrix} | & | & | & | & | \\ \tilde{\mathbf{x}}_1 & \tilde{\mathbf{x}}_2 & \tilde{\mathbf{x}}_3 & \dots & \tilde{\mathbf{x}}_m \\ | & | & | & | & | \end{bmatrix} \quad (5)$$

U consists of orthogonal vectors representing the modes of variability. The Σ matrix contains the eigenvalues—corresponding to the columns in U —along the diagonal, and V^T is composed of the right singular vectors of \mathbf{X} . The data is encoded by performing matrix multiplication with U .

Neural network modeling. In this work, we leverage neural networks (NNs) for nonlinear regression modeling due to their applicability as universal function approximators and flexibility in development. A neural network is a collection of computational cells (or neurons) connected in some form through multiplicative connections (or weights). Artificial neural networks (ANNs) have been used to directly predict thermospheric mass density using space weather indices and proxies as model drivers in order to study long-term trends^{55,56}. These types of models can also be used as an exercise in understanding the effect of the drivers on non-machine learning (ML) models⁵⁷. Chen et al.⁵⁸ developed ANNs with different combinations of geomagnetic indices to fit to CHAMP and GRACE density estimates during storms, and Choury et al.⁵⁹ developed an ANN to predict exospheric temperature for use in the Drag Temperature Model (DTM).

Loss functions. Loss functions are used to inform the NN of the objective during the training phase, or weight adjustment period. Loss functions can be minimized or maximized depending on the modeling objective. In this work, we minimize the negative logarithm of predictive density (NLPD) given as,

$$NLPD(y, \mu, \sigma) = \frac{1}{n} \sum_{i=1}^n \left(\frac{(y_i - \mu_i)^2}{2\sigma_i^2} + \frac{\ln(\sigma_i^2)}{2} + \frac{\ln(2\pi)}{2} \right) \quad (6)$$

where y is the observed value, μ is the predicted mean, σ is the standard deviation of the output corresponding to each unique input, and n is the batch size. The batch size is the number of samples the model will pass through before updating the weights, averaging the loss over the batch. Losses are computed for every output

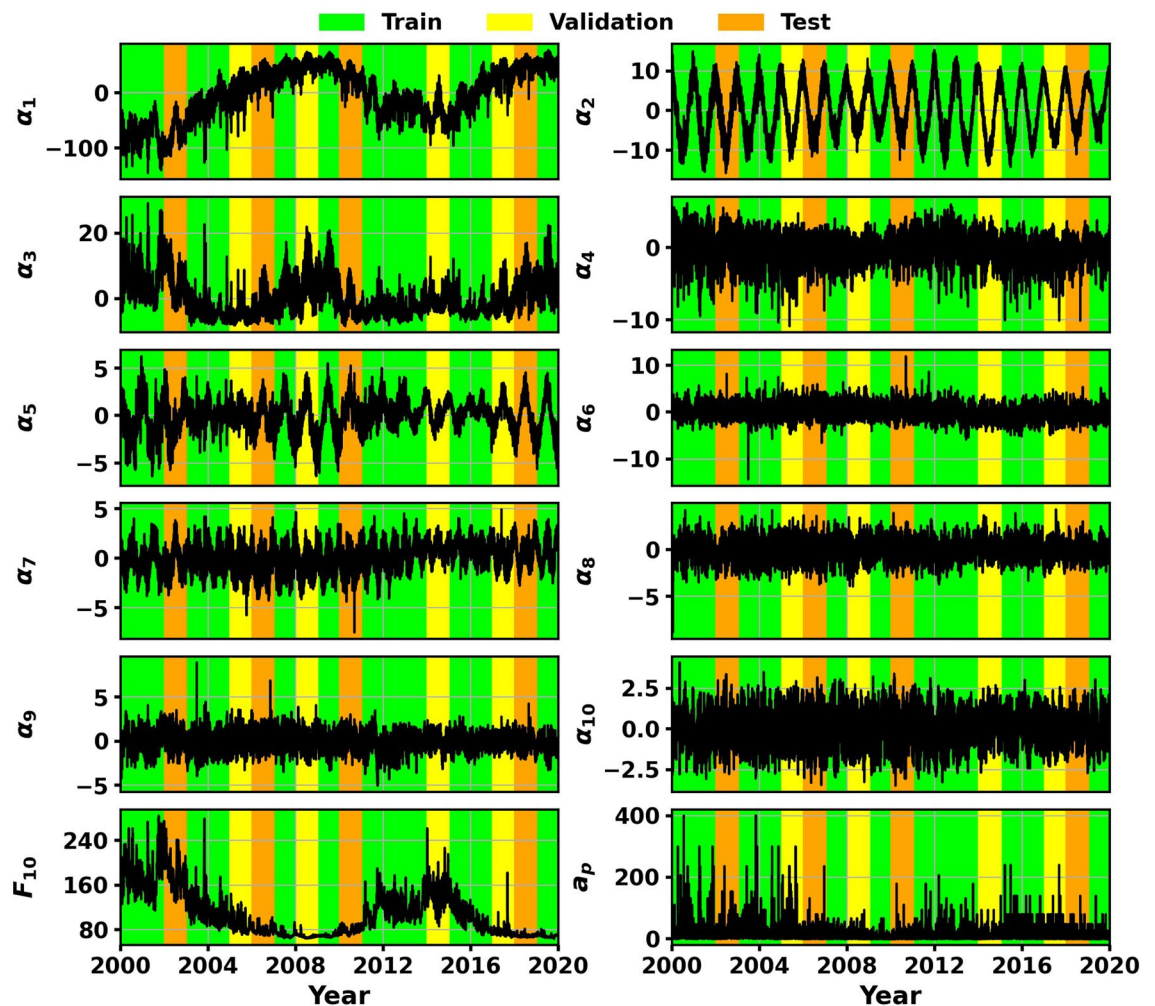


Figure 2. First 10 HASDM PCA coefficients with F_{10} and a_p . The shading represents the training, validation, and test sets.

and backpropagation is how the model determines how much to change each weight^{60,61}. NLPD is derived from $-\ln(f(x))$ where \ln is the natural logarithm and $f(x)$ is the probability density function of the normal distribution.

Hyperparameter optimization. Tools like Keras Tuner have drastically reduced model development time⁶². You can provide ranges of hyperparameters and Keras Tuner will explore the search space. We use the Bayesian optimization scheme, allowing the tuner to perform a random search for the first 25 trials, or architectures, and using a Gaussian process (GP) model to choose the architectures for the final 75 trials to exploit the high performing areas of the space. The objective of the tuner is to minimize validation loss. The model optimizer and number of layers are first chosen by the tuner. For each layer, the model can have a unique number of neurons, activation function, and dropout rate. For each model developed (two datasets and two UQ techniques), the architecture is selected using Keras Tuner. It is important to note that the tuner can provide multiple architectures that provide similar results.

The 58,440 samples in the HASDM dataset are split into 60% training, 20% validation, and 20% test data. This is displayed in Fig. 2. As the number of training and validation samples is manageable, the full sets are used in tuning. We obtain the HASDM models directly from the tuner without a need to train further.

The CHAMP dataset is significantly larger with over 25 million total samples. Unlike the HASDM dataset, location is now an input. CHAMP only covers the local solar time domain once every 3 months. The dataset also does not span an entire solar cycle. We originally tried using a 10-month segment of 2003 for validation and a 10-month segment from 2005 to 2006 for testing. This resulted in poor model generalization due to the lack of coverage of the solar cycle in the remaining training set. Therefore, we repeated the following data split scheme. Eight weeks are used for training (483,840 samples), then the following week is used for validation (60,480 samples), and the next week is used for the test set (60,480 samples). This results in similar input and output distributions while keeping temporally disjoint sets as there are 2 weeks or 120,960 samples between the training segments. For the tuner, 1 million random samples are chosen from the training data and 500,000 random samples are chosen from the validation data. Once the tuner is complete, the best models are retrained on the full training set and evaluated on the other two sets.

	Function	σ
Problem 1	$0.3x + \cos(0.5x) - 4 + \mathcal{N}(0, \sigma)$	$0.5 \frac{e^{\sin(0.2x)}}{1 + e^{\sin(0.8x)}}$
Problem 2	$\sin(2x + \cos(3x)) + \mathcal{N}(0, \sigma)$	$0.05\sin(0.2x)$

Table 2. Functions for the two toy problems with the right column being the functional form of the Gaussian noise.

Uncertainty quantification. We use two ML techniques: MC dropout and direct probability distribution prediction, as UQ with machine-learned models is fairly unexplored in the space weather domain. Dropout is a generalization technique that applies Bernoulli distributions in each layer to change the flow of information through the model^{63,64}. Dropout is traditionally only active during training to maintain a deterministic form in prediction. By forcing dropout to remain active in prediction, the model becomes probabilistic. MC dropout has been shown to be an approximation of a GP⁶⁵. For both methods, we use the negative logarithm of predictive density (NLPD) loss function (Eq. 6). Licata et al. found that the mean square error loss function resulted in underestimated uncertainty estimates in surrogate modeling for the HASDM dataset²³.

Monte Carlo dropout implementation. The typical input and output shape is $n \times n_{\text{inp}}$ and $n \times n_{\text{out}}$, respectively. n is the number of samples, n_{inp} is the number of inputs, and n_{out} is the number of outputs. In training, the mean and standard deviation need to be unique to each input sample, so the model has to be provided each input k times. k needs to be a large enough number to allow for adequate representation of the predicted distribution. The inputs and outputs for training are stacked about a repeated intermediary axis. The training samples are identical about k , but are unique about n . The new input and output shapes—necessary for proper training—are $n \times k \times n_{\text{inp}}$ and $n \times k \times n_{\text{out}}$, respectively. In each training batch, the mean and standard deviation are taken with respect to the intermediate axis, and the NLPD loss can be computed.

Direct probability distribution prediction. Another way to represent uncertainty is to directly predict the mean and standard deviation. The mean square error loss function cannot be used here as there are no labels for the standard deviation. However, Nix and Weigend used a neural network to directly predict the mean and variance of a toy dataset using the NLPD loss function⁶⁶. We implement this technique for the datasets presented. To accomplish this, we create a custom output layer with $2n_{\text{out}}$ neurons. The first n_{out} neurons represent the mean prediction and have a linear activation function. The last n_{out} neurons represent the standard deviation and use the softplus activation function. The softplus function and its derivative—the sigmoid function—are shown in Eq. (7).

$$f(x) = \ln(1 + e^x) \quad f'(x) = \frac{e^x}{1 + e^x} \quad (7)$$

The desired qualities of the standard deviation output are: (1) always positive and (2) having no upper bound. The initial choice was the absolute value function. However, the resulting models had erratic loss values, and it was difficult to obtain a good model. The softplus function is (1) always positive, (2) has no upper bound, (3) is monotonically increasing, and (4) is differentiable across all inputs. This resulted in stable training losses and better models.

Metrics. To compare the predictive capability of the models developed, we look at the mean absolute error across the training, validation, and test sets. The errors across different space weather conditions will be investigated as well. We also test the reliability of the uncertainty estimates both qualitatively and quantitatively. The calibration error score is given as

$$\text{Calibration Error} = \frac{100\%}{m \cdot n_{\text{out}}} \sum_{i=1}^{n_{\text{out}}} \sum_{j=1}^m \left| p(z_{i,j}) - p(\hat{z}_{i,j}) \right| \quad (8)$$

where m is the number of prediction intervals (PIs) of interest. Here the PIs range from 5 to 95% with 5% increments in addition to 99%—[0.05, 0.10, 0.15, ..., 0.90, 0.95, 0.99]. $p(\hat{z}_{i,j})$ is the observed cumulative probability obtained by dividing the number of true samples within the prediction interval by the total number of samples. Equation (8) is the miscalibration of prediction intervals averaged over each output and prediction interval tested. For this work, it provides the average deviation from all 20 PIs for each model output. We can visualize the reliability of the uncertainty estimates by plotting the calibration curve— $p(\hat{z})$ vs $p(z)$.

Toy problems

To visualize the way the NLPD loss function influences training, we train models for two toy problems. Each problem is a function, $y(x)$, with additive Gaussian noise having zero-mean and a functional form to the standard deviation. These functions are displayed in Table 2. The results for Problem 1 is shown in Fig. 3.

Figure 3 shows that the model is able to adequately predict the function and is able to predict the overall probability distribution. The interesting aspect of the figure is panel (d): the model is able to predict the standard

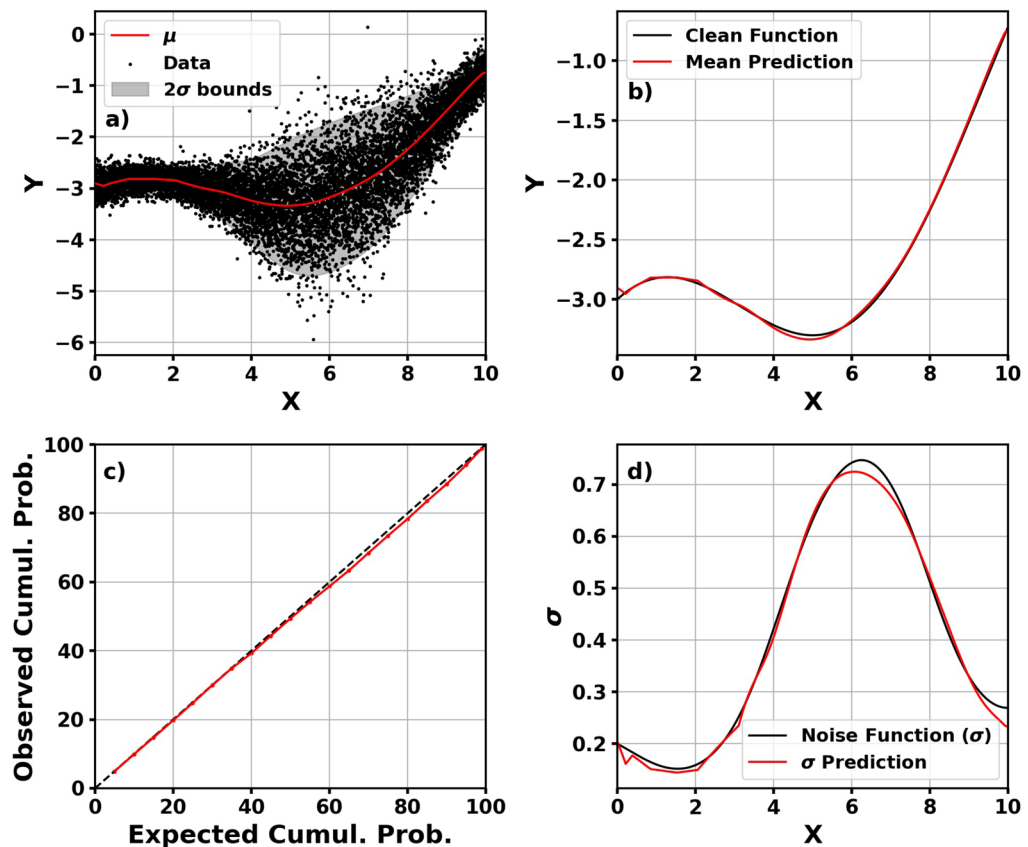


Figure 3. Mean prediction with 2σ bounds plotted on data (a), clean function plotted with mean prediction (b), calibration curve (c), and predicted standard deviation on true standard deviation function (d) for Problem 1.

deviation without a label. Meanwhile, this is fairly trivial data. Figure 4 shows the predictions and calibration curve for the more complex Problem 2.

For the more complex data, the model is not as accurate overall x . When $x < 6$, the model can accurately predict the mean and standard deviation. When $x > 6$, the standard deviation prediction no longer represents uncertainty in the data but the model's uncertainty in its prediction. For this portion of panel (b), the mean prediction deviates from the true mean of the data and the standard deviation in panel (d) consequently increases. Panel (c) shows that the model is still well-calibrated and representing both uncertainty in the data and uncertainty in the model's predictions.

The NLPD loss function does not ensure model calibration. However, we show that it can be used—if properly tested—in model development to represent uncertainty in the data and uncertainty in the model's predictions. Note: these models were trained on the entire dataset, and this is purely for demonstration. The thermospheric density models are developed with separate validation and independent test sets.

HASDM model

Using the best tuner models for MC dropout and direct probability distribution prediction, we assess the error and calibration statistics. Table 3 shows the mean absolute error and calibration error score for both techniques across the training, validation, and test sets.

It is evident that the performance using both methods is very similar. Across all three sets, the mean absolute error and calibration error score do not deviate by more than 0.8% and 1.4% respectively. The MC dropout model has better performance on the independent test set in terms of calibration. This is a desired quality as the test data is not used for model development in any way. As the calibration error scores are composites of the scores for each output, the calibration curves are shown in Fig. 5 for a qualitative assessment.

Both techniques lead to slightly overestimated uncertainties on the training set for multiple outputs. Meanwhile, the remaining outputs are almost perfectly calibrated. On the validation set, each model has outputs with overestimated and underestimated uncertainties. Again, most of the outputs are very well-calibrated which is affirmed by the calibration error scores. For the test set, the direct probability prediction model tends to marginally underestimate the uncertainty while the MC dropout model provides reliable uncertainty estimates on virtually all model outputs. Table 4 shows the mean absolute error for both models across an array of solar and geomagnetic conditions. The entire dataset is used for this analysis as there are not enough samples in each bin using only the test set.

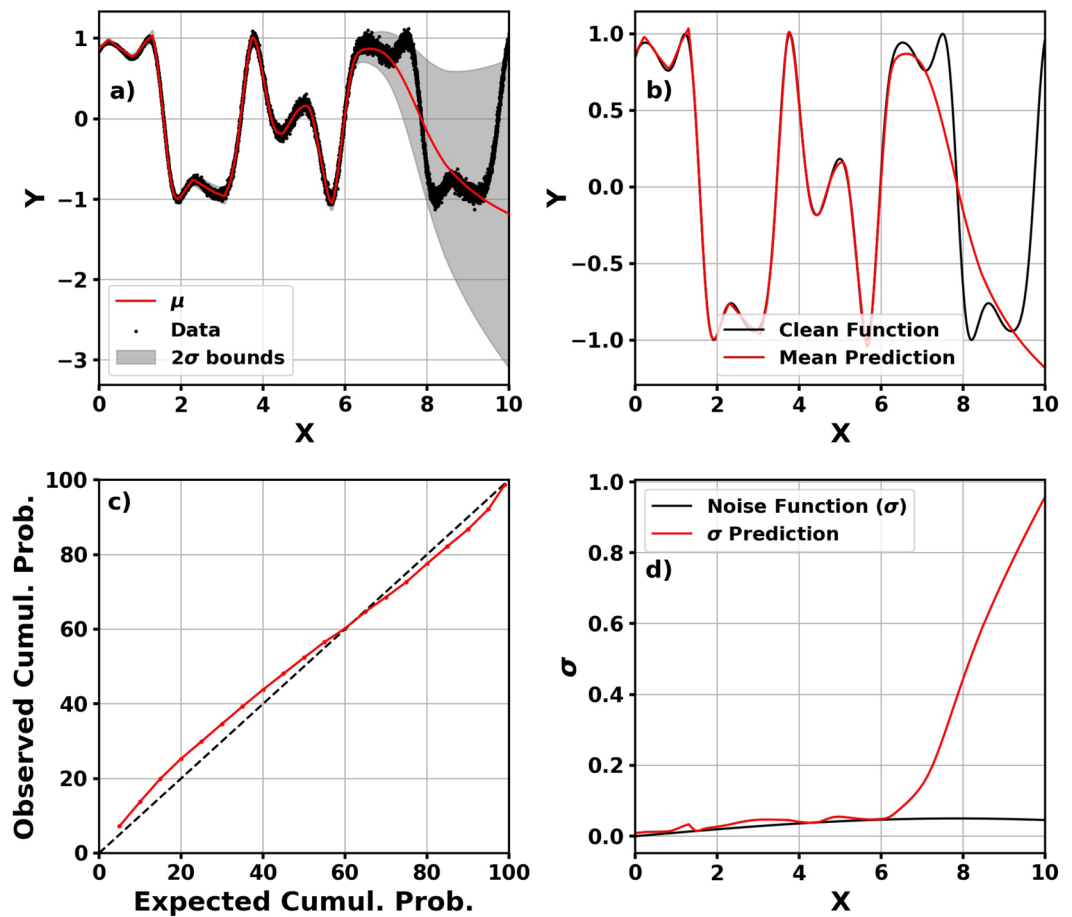


Figure 4. Mean prediction with 2σ bounds plotted on data (a), clean function plotted with mean prediction (b), calibration curve (c), and predicted standard deviation on true standard deviation function (d) for Problem 2.

Metric	Set	MC dropout (%)	Direct probability (%)
Error	Training	9.07	8.55
	Validation	10.69	9.91
	Test	10.69	10.60
Calibration	Training	3.06	1.74
	Validation	2.51	2.45
	Test	1.76	2.81

Table 3. HASDM modeling results using MC dropout and direct probability prediction. Error refers to mean absolute error, and calibration is computed using Eq. (8). Bold values represent the lowest error between methods.

These errors tend to reiterate the results from Table 3. The direct probability model was more accurate on all three sets, and Table 4 shows that it is also more accurate across all 20 conditions considered. For a majority of the conditions, the difference is small ($< 1\%$). However, the high a_p conditions show that the direct probability model makes considerable improvements. These error reductions from MC dropout range from 1.6 to 4.1%.

To further assess the uncertainty capabilities of the models, we attempt to visualize the calibration in the full-state (global density grids) to identify any spatial dependence in the reliability of the uncertainty estimates. First, the models are evaluated on the entire test set and the density mean and standard deviations are extracted. Using these statistics, the observed cumulative probability with a 90% prediction interval is computed for each spatial location. The resulting $24 \times 19 \times 27$ array is used to determine how well calibrated the model is on independent data as a function of location. We show seven maps for each model (200, 300, ..., 800 km) in Fig. 6. Even though HASDM has a lateral spatial resolution of 24 longitude and 19 latitude segments, we interpolate the results to the polyhedral grid used in the EXEMPLAR model for visualization purposes. This is done in the remainder of the manuscript.

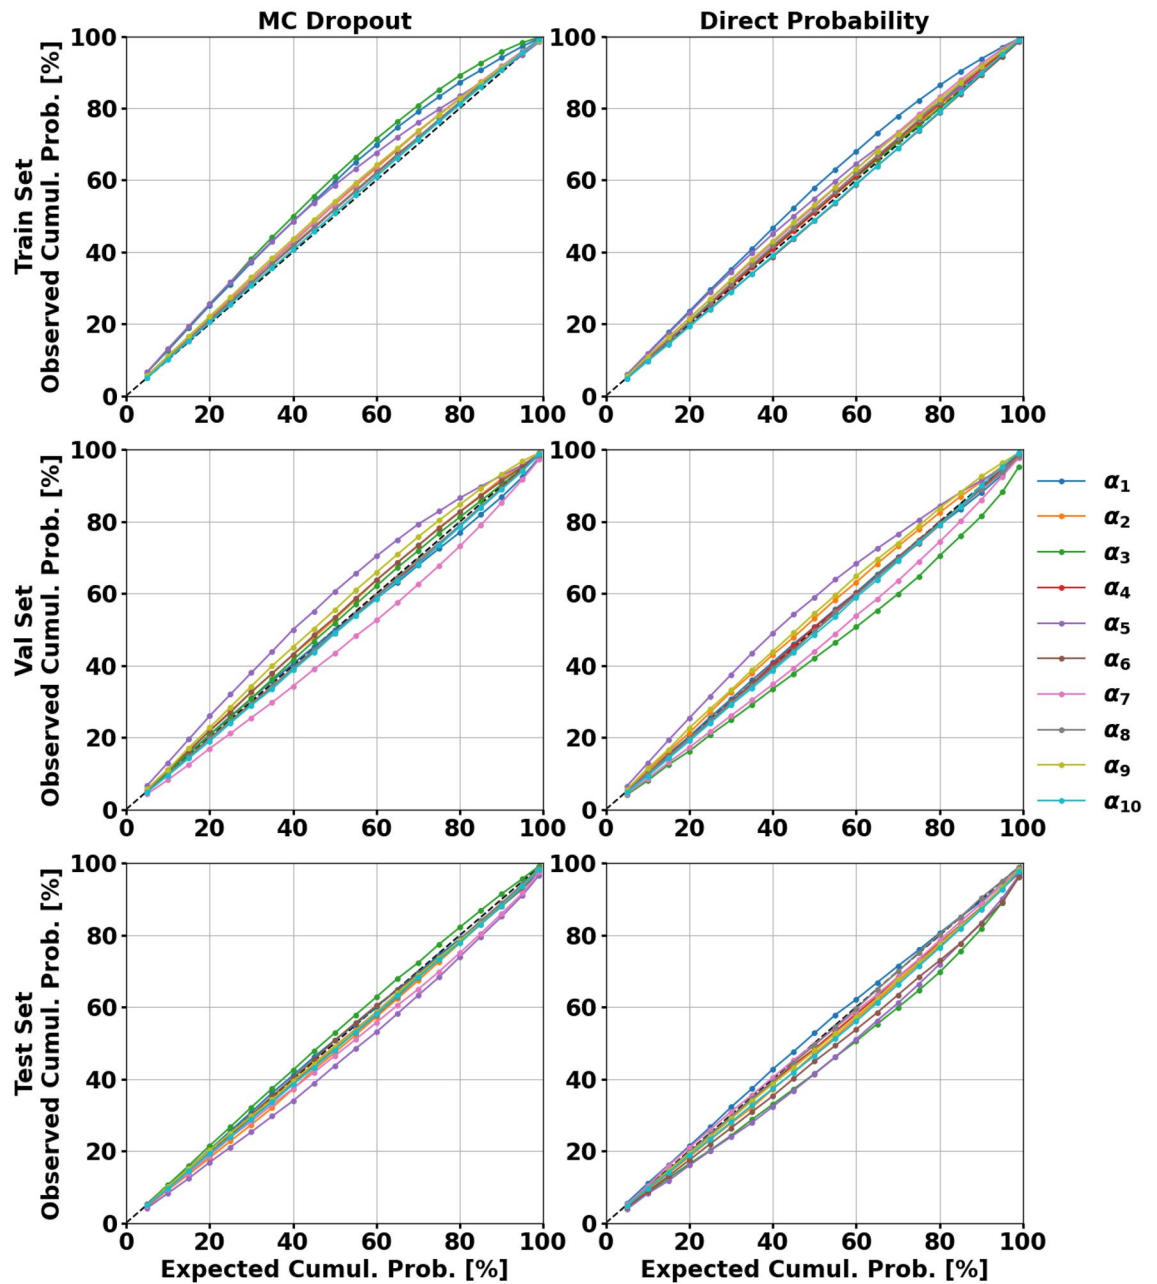


Figure 5. The left and right columns show the MC dropout and direct probability calibration curves, respectively. The top, middle, and bottom rows are the calibration curves for the training, validation, and test sets, respectively.

For reference, perfect calibration in Fig. 6 would be uniform green maps at all altitudes. This would convey that with a 90% prediction interval, the model's predictions/uncertainty estimates contain 90% of true samples at all locations. While this is not the case, the results are still insightful. At 200 km, both models are underestimating the uncertainty by 10–15%. This could be a result of the relative variability as a function of altitude in the SET HASDM density database. The general trend of relative variability is that it increases with altitude, so the models may underpredict the standard deviation at low altitudes as a result, which indicates that the model has a false sense of confidence in that region. Both models have an average cumulative probability within 5% of the expected value at most of the altitudes shown in Fig. 6 with the best results at 600 km. At 700 and 800 km, both models begin to overestimate uncertainty, likely because they have the lowest confidence at those altitudes. An interesting outcome of this study is the lateral variability of the cumulative probability between the models. The MC dropout model (left) has more lateral variability, meaning the cumulative probability changes more as a function of longitude and latitude.

	$F_{10} \leq 75$	$75 < F_{10} \leq 150$	$150 < F_{10} \leq 190$	$F_{10} > 190$	All F_{10}
MC dropout					
$a_p \leq 10$	8.96%	9.78%	9.97%	9.14%	9.50%
$10 < a_p \leq 50$	9.76%	10.05%	10.87%	9.90%	10.09%
$a_p > 50$	15.35%	12.86%	13.23%	12.55%	13.01%
All a_p	9.12%	9.92%	10.36%	9.55%	9.71%
Direct probability					
$a_p \leq 10$	8.64%	9.33%	9.35%	9.11%	9.10%
$10 < a_p \leq 50$	9.18%	9.51%	9.69%	9.64%	9.48%
$a_p > 50$	11.14%	11.23%	11.34%	10.30%	11.11%
All a_p	8.74%	9.42%	9.52%	9.34%	9.23%

Table 4. Mean absolute error across global grid for HASDM-ML as a function of space weather conditions.




Figure 6. Observed cumulative probability maps for a 90% prediction interval using the MC dropout (left) and direct probability (right) models. The average observed cumulative probability is shown for each altitude in parenthesis.

CHAMP model

After running tuners for both uncertainty techniques, we trained the best models on the entire training set. The models were chosen based on the lowest prediction error and best calibration scores on the validation set. Table 5 shows the mean absolute error and calibration error scores on the three sets.

Both models are well-generalized in terms of prediction accuracy. The range in error between sets for the MC dropout and direct probability model is 0.54% and 0.23%, respectively. Both models have higher calibration

Metric	Set	MC dropout (%)	Direct probability (%)
Error	Training	13.13	12.59
	Validation	13.67	12.82
	Test	13.14	12.62
Calibration	Training	3.93	5.84
	Validation	0.64	0.25
	Test	0.22	0.37

Table 5. CHAMP modeling results using MC dropout and direct probability prediction. Error refers to mean absolute error, and calibration is computed using Eq. (8). Bold values represent the lowest error between methods.

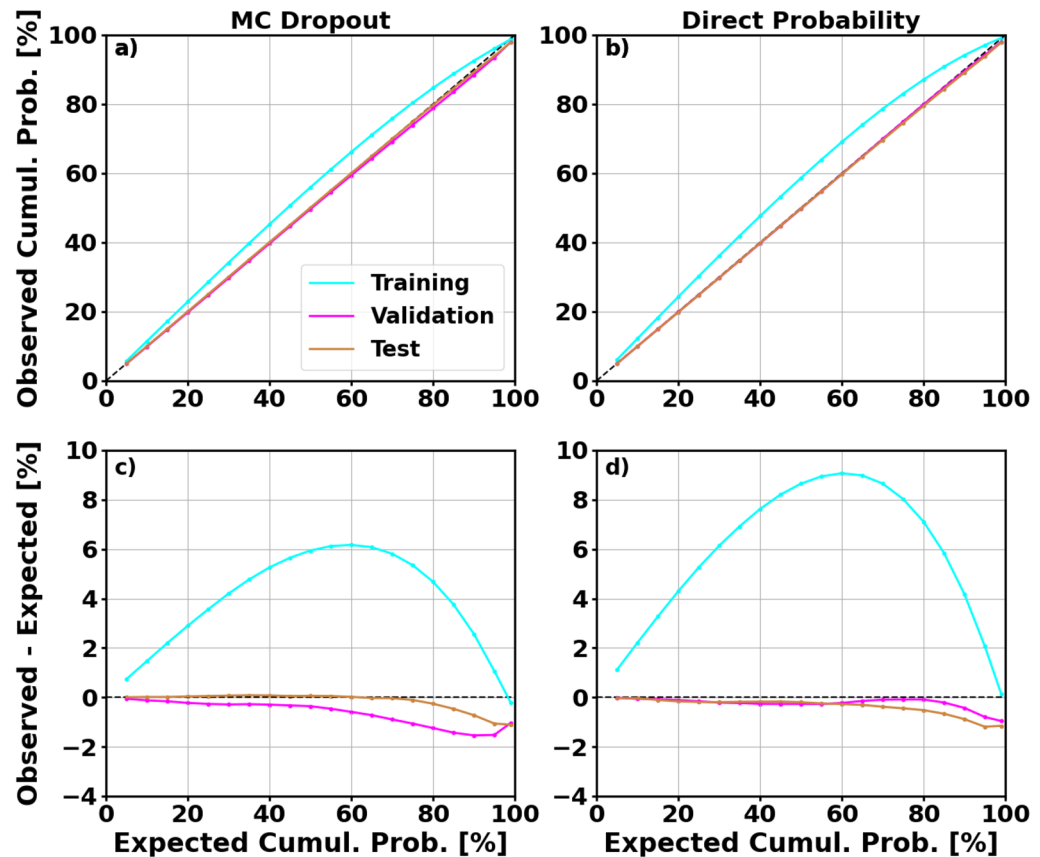


Figure 7. Calibration curves for the training, validation, and test sets using MC dropout (a) and direct probability prediction (b). (c) and (d) show the difference between the observed and expected cumulative probability using MC dropout and direct probability prediction, respectively.

error scores on the training set but have similar scores on the validation and test sets. The two techniques provide similar results with the only notable difference is the 1.91% higher calibration error score for the direct probability model on the training set. The calibration curves for both models are shown in Fig. 7.

Both models are well-calibrated on all three sets. There is a tendency for both models to slightly overestimate uncertainty on the training set which is more evident for the direct probability model. The differences between the calibration curves and the perfectly calibrated reference line (in black) is shown in panels (c) and (d). Panel (d) highlights the overestimation of uncertainty for the direct probability model on the training set. However, it never deviates by more than 9%. Both models tend to underestimate uncertainty on the validation and test set for the larger prediction intervals. Again, the deviation from perfect calibration is no more than 2% for any PI. Due to the intrinsic difference between the datasets that the CHAMP and HASDM models are developed from, the proceeding analyses will be different than those for the HASDM model.

Global modeling with local measurements. The CHAMP models were developed with in-situ measurements, but we hypothesize that it should be able to learn the functional relationship of the combined inputs. Therefore, the model should be able to provide global outputs at any point in time. As a qualitative assessment,

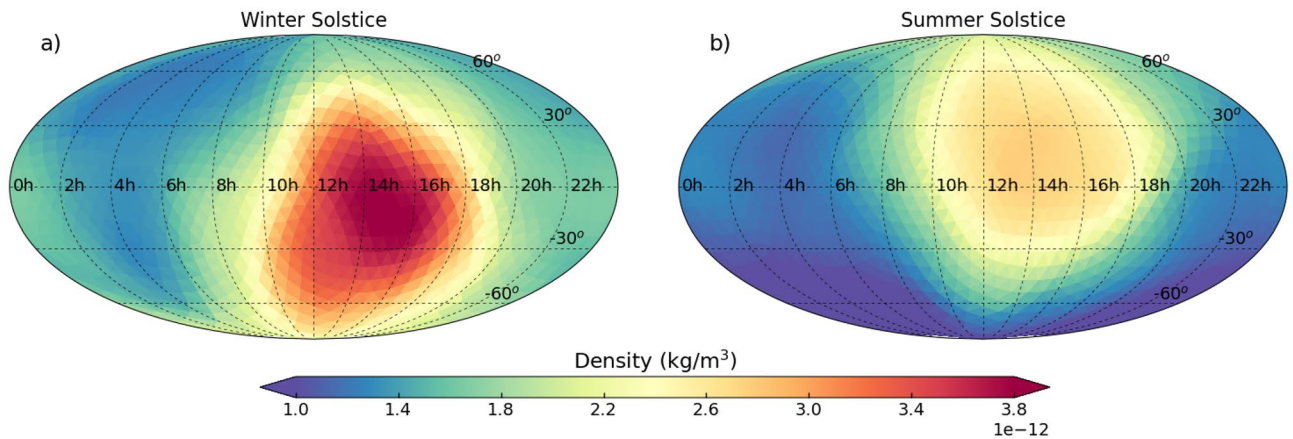


Figure 8. Global density map with moderate solar activity, low geomagnetic activity, the altitude fixed to 400 km, and the time of day being 00:00 UTC for the winter solstice (a) and the summer solstice (b).

Condition name	Solar drivers	Geomagnetic drivers		Temporal drivers	
	FMSY	SYM-H	$S_N = S_S$	UTC	doy
Solar 1	75	0	27	0	262
Solar 2 ^a	120	0	27	0	262
Solar 3	190	0	27	0	262
Geo 2	120	-75	128	0	262
Geo 3	190	-75	128	0	262
UTC 2	120	0	27	12	262
doy 2	120	0	27	0	172
doy 3	120	0	27	0	355

Table 6. CHAMP model inputs to study various conditions as a function of altitude. ^aSolar 2 is also considered Geo 1, UTC 1, and doy 1.

we show global maps at 400 km for the winter and summer solstices in Fig. 8 using the direct probability model. All proceeding global analyses will be performed using this model. For this test, the solar drivers are all set to 120 sfu, *SYM-H* is set to 0 nT, both Poynting flux totals are set to 27GW, and the time is set to 00:00 UTC.

The diurnal structure is present in both panels with the peak density being in the southern hemisphere during the winter solstice and in the northern hemisphere during the summer solstice. This shows the model's understanding on annual trends (Earth's tilt). The general density level is higher during the winter solstice, but the relative variation between day and night are very similar. This is reaffirmed by the exospheric temperature distribution shown by Weimer et al.⁴³ during the solstices. Additional global density maps at different altitudes can be found in Fig. S1 using baseline conditions. Furthermore, a global storm example is shown in Figs. S2 and S3.

Next, we look at the uncertainty levels for eight unique conditions of activity and time. These are all displayed in Table 6. Using these space weather and temporal inputs, the CHAMP model is evaluated at all 1620 polyhedral grid locations from 300 to 450 km in 1 km increments. The metric we use here is a normalized measure of model uncertainty: $100 \cdot \sigma / \mu$, essentially providing the $1 - \sigma$ uncertainty as a percentage of the mean prediction. The resulting maps are averaged across each altitude to evaluate the model's uncertainty for each condition as a function of altitude. Three aspects of model drivers are investigated: solar activity, geomagnetic activity, and temporal dependence. In Table 6, there are three solar activity levels, with all other drivers kept constant. There are also three geomagnetic cases: low and high geomagnetic activity with moderate solar activity, and high geomagnetic activity with high solar activity. We only look at two daily cases—00:00 and 12:00 UTC. We also look at the fall equinox, summer solstice, and winter solstice with moderate solar and low geomagnetic activity. The resulting altitude profiles are shown in Fig. 9.

Panel (a) in Fig. 9 shows that the CHAMP model has low uncertainty in its lower altitude predictions for solar minimum (or low solar activity) which drastically increases with altitude. The opposite can be said for solar maximum. The moderate solar activity case results in lowest uncertainties between 350 and 375 km and higher uncertainties above and below that range. This is all a result of CHAMP's altitude from 2002 to 2010. It started around 460 km during solar maximum and ended at 300 km during solar minimum. Therefore, the model has confident predictions in the altitude range the satellite was located during the various phases of the solar cycle. If there was additional data from satellites at different altitudes over a longer time period, the model would likely be more confident over a larger altitude range.




Figure 9. Normalized uncertainty variations as a function of altitude for solar (a), geomagnetic (b), daily (c), and annual (d) cases. The drivers for each curve can be found in Table 6.

In panel (b), we see the same general trends for Geo 1 and Geo 2, because they are evaluated using moderate solar activity. However, it is evident that the increase in geomagnetic activity results in up to 5% more uncertainty. The Geo 3 case is similar to Solar 3 (high solar activity) but again has increased uncertainty due to the storm conditions it represents. Panel (c) indicates that there is a low impact from universal time on the model uncertainty. In Panel (d), the black line indicates the fall equinox which is similar to the winter solstice. The Winter solstice uncertainties deviate from the equinox uncertainties at the highest altitude range. While the overall shape remains consistent, there are highest uncertainties for the summer solstice at all altitudes. The overall takeaway from Fig. 9 is that the shape of the model uncertainty altitude profile is most strongly effected by the solar activity level while the day of year and geomagnetic activity tend to uniformly increase or decrease uncertainty. These profiles would all likely be impacted if the model was developed using additional satellite data.

Evaluation time comparison

We attempt to provide an equal comparison of the two methods in terms of computational complexity. To do so, each CHAMP model is evaluated on either 8640 samples (1 week) or 86,400 samples (10 weeks). For the direct probability prediction model, it sees each input once and provides the mean and standard deviation. These are used to sample a Gaussian distribution 1000 times to get probabilistic predictions for density over the given window.

For MC dropout, we cannot pass 1 week of inputs to the model stacked 1000 times (as is done for HASDM). There is not enough memory on an NVIDIA GeForce RTX 2080 Ti graphics processing unit (GPU)—11 GB—to perform this evaluation. Therefore, we pass the 100 repeated inputs in 10 chunks to obtain the 1000 predictions. When evaluating over 10 weeks, we must reduce to 10 repeated inputs in 100 chunks. In Table 7, we show the evaluation times on both GPU and CPU for both methods over the two durations. Note: when running MC dropout on CPU, we use 100 repeated inputs for both durations. The batch size for all predictions is 2^{17} or 131,072. The size of the MC dropout and direct probability models are 233.3 kB and 21.9 MB, respectively.

The run times are unique to these specific models. The size of the models plays a role in run time, and the size of these models are a result of the tuner. The MC dropout model is approximately 100 times smaller, but the increase in required model prediction calls results in the significantly longer run times. The direct probability

Method	Samples	GPU run time	CPU run time
MC dropout	8640	2.11	13.65
	86,400	18.29	127.79
Direct probability	8640	0.58	0.52
	86,400	3.93	3.93

Table 7. Run time to obtain 1000 probabilistic predictions from each model using GPU and CPU in seconds.

Method	Pros	Cons
MC dropout	No need to sample from a Gaussian distribution	Longer evaluation times
		Not compatible with large datasets
Direct probability	Only need single evaluation	Required sampling from a Gaussian distribution to obtain probabilistic predictions
	Computationally efficient	

Table 8. Pros and cons for MC dropout and direct probability distribution prediction.

method, for this particular problem, is anywhere from 3 to 30 times faster depending on the number of samples and whether the GPU or CPU is being used.

Discussion

In this work, we leverage the NLPD loss function to develop thermospheric density models using (1) MC dropout and (2) direct probability prediction. These two uncertainty techniques were used to create both a model based in the PCA coefficients of the SET HASDM density database and a model based in localized accelerometer-derived density estimates from CHAMP. Using two toy problems, we showed that the NLPD loss function can be used to create a ML model with calibrated uncertainty estimates relative to uncertainty in the model, uncertainty in the data, or both. For the HASDM database, the MC dropout and direct probability distribution prediction models had similar metrics in terms of error and calibration. Furthermore, the calibration curves for the PCA coefficients were nearly identical. By looking at the density calibration of both HASDM models, we found that they were well-calibrated at mid-altitudes, and there was more lateral variability in the calibration of the MC dropout model.

The CHAMP models also had similar performance and were both well-generalized. We test the CHAMP model's global prediction capabilities by generating baseline maps during the winter and summer solstices to ensure physical global trends are being captured by the CHAMP model. This showed that the model was able to emulate the effect of Earth's tilt. We also performed global evaluations for eight unique conditions to determine the altitude dependence of model uncertainty. The altitude profiles showed that the minimum and maximum $1 - \sigma$ uncertainties were 10–28% of the mean predictions, respectively. Solar activity was most influential in determining the profiles' shapes while geomagnetic activity and day of year tended to provide uniform changes in the uncertainty. In general, the MC dropout and direct probability methods were shown to have similar performance for thermospheric density modeling applications. However, there are pros and cons for both methods, and careful consideration is required when deciding on a UQ method for space weather models. These are highlighted in Table 8.

The main disadvantage for the direct probability method is the requirement to sample from a Gaussian distribution to get probabilistic density variations. The drawback to MC dropout is its higher computational cost. In terms of density modeling, both techniques have prompt evaluation times. Relative to one another, we show that the direct probability models can be evaluated much faster. The size of the training data (in both number of samples and dimensionality) is also important to consider. With MC dropout, GPU memory can constrain modeling efforts if the dataset is too large. It can also require additional steps for prediction. In this work, MC dropout did not inhibit model development for the smaller HASDM PCA data. However, it did add numerous considerations when developing and evaluating the CHAMP model. In general, the uncertainty estimation capabilities may be improved through modifications to the loss function to either: (a) add higher order moments or (b) obtain non-Gaussian estimates.

All the preceding results show that for thermospheric density applications, these two techniques can be used to obtain an accurate model with reliable uncertainty estimates. There are other methods that can be used in space weather application such as GP regression and ensemble modeling, but this is a sufficient starting point. Other final considerations concern orthogonality and applicability. For a multi output regression model (e.g. HASDM models), the outputs must be orthogonal. This is to both prevent collinearity and since the use of NLPD requires uncorrelated outputs. The CHAMP data only spans an altitude range of 300–460 km. Any predictions outside of this range may be unreliable. To combat this, density estimates from other satellites can be added to increase the altitude coverage and provide the model with more data to learn from, as discussed in “[Satellite accelerometer density estimates](#)”.

Data availability

Requests can be submitted for access to the SET HASDM density database at <https://spacewx.com/hasdm/> and all reasonable requests for scientific research will be accepted as explained in the rules of road document on the website. The historical space weather indices used in this study can also be found at <https://spacewx.com/jb2008/> in the SOLFSMY.TXT, SOLRESAP.TXT, and DSTFILE.TXT files for the solar indices, *ap*, and *Dst*, respectively. Free and one-time only registration is required to access these files. The forecasting capabilities for the solar drivers has been recently benchmarked⁶⁷. *SYM-H* data was obtained from <http://wdc.kugi.kyoto-u.ac.jp/aeasy/index.html> thanks to the World Data Center for Geomagnetism in Kyoto. CHAMP density estimates from Mehta et al.²⁹ can be found at <http://tinyurl.com/densitysets>.

Received: 12 January 2022; Accepted: 18 April 2022

Published online: 04 May 2022

References

- Radtke, J., Kebschull, C. & Stoll, E. Interactions of the space debris environment with mega constellations—Using the example of the OneWeb constellation. *Acta Astronautica* **131**, 55–68. <https://doi.org/10.1016/j.actaastro.2016.11.021> (2017).
- Pardini, C. & Anselmo, L. Revisiting the collision risk with cataloged objects for the Iridium and COSMO-SkyMed satellite constellations. *Acta Astronautica* **134**, 23–32. <https://doi.org/10.1016/j.actaastro.2017.01.046> (2017).
- Boley, A. C. & Byers, M. Satellite mega-constellations create risks in Low Earth Orbit, the atmosphere and on Earth. *Sci. Rep.* **11**, 1–8. <https://doi.org/10.1038/s41598-021-89909-7> (2021).
- Roble, R. G. *Energetics of the Mesosphere and Thermosphere* 1–21 (American Geophysical Union (AGU), 1995).
- Qian, L. & Solomon, S. Thermospheric density: An overview of temporal and spatial variations. *Space Sci. Rev.* **168**, 1–27. <https://doi.org/10.1007/s11214-011-9810-z> (2011).
- Tobiska, W. K., Bouwer, S. D. & Bowman, B. R. The development of new solar indices for use in thermospheric density modeling. *J. Atmos. Solar-Terr. Phys.* **70**, 803–819. <https://doi.org/10.1016/j.jastp.2007.11.001> (2008).
- Bruinsma, S. L. & Forbes, J. M. Properties of traveling atmospheric disturbances (TADs) inferred from CHAMP accelerometer observations. *Adv. Space Res.* **43**, 369–376. <https://doi.org/10.1016/j.asr.2008.10.031> (2009).
- Oliveira, D. M., Zesta, E., Schuck, P. W. & Sutton, E. K. Thermosphere global time response to geomagnetic storms caused by coronal mass ejections. *J. Geophys. Res. Space Phys.* **122**, 10762–10782. <https://doi.org/10.1002/2017JA024006> (2017).
- Bruinsma, S., Boniface, C., Sutton, E. K. & Fedrizzi, M. Thermosphere modeling capabilities assessment: Geomagnetic storms. *J. Space Weather Space Clim.* **11**, 12. <https://doi.org/10.1051/swsc/2021002> (2021).
- Ritter, P., Lühr, H. & Doornbos, E. Substorm-related thermospheric density and wind disturbances derived from CHAMP observations. *Ann. Geophys.* **28**, 1207–1220. <https://doi.org/10.5194/angeo-28-1207-2010> (2010).
- Bruinsma, S. & Biancale, R. Total densities derived from accelerometer data. *J. Spacecr. Rockets* **40**, 230–236. <https://doi.org/10.2514/2.3937> (2003).
- Liu, H., Lühr, H., Henize, V. & Köhler, W. Global distribution of the thermospheric total mass density derived from CHAMP. *J. Geophys. Res. Space Phys.* <https://doi.org/10.1029/2004JA010741> (2005).
- Sutton, E. K. *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 (2008). <https://ui.adsabs.harvard.edu/abs/2008PhDT.....87S>.
- Doornbos, E. *Producing Density and Crosswind Data from Satellite Dynamics Observations* 91–126 (Springer, 2012).
- Calabia, A. & Jin, S. New modes and mechanisms of thermospheric mass density variations from GRACE accelerometers. *J. Geophys. Res. Space Phys.* **121**, 11191–11212. <https://doi.org/10.1002/2016JA022594> (2016).
- Picone, J. M., Hedin, A. E., Drob, D. P. & Aikin, A. C. NRLMSISE-00 empirical model of the atmosphere: Statistical comparisons and scientific issues. *J. Geophys. Res. Space Phys.* **107**, S1A 15-1–S1A 15-16. <https://doi.org/10.1029/2002JA009430> (2002).
- Bowman, B. et al. A new empirical thermospheric density model JB2008 using new solar and geomagnetic indices. In *AIAA/AAS Astrodynamics Specialist Conference* (AIAA 2008-6438, 2008). <https://doi.org/10.2514/6.2008-6438>.
- Bruinsma, Sean. The DTM-2013 thermosphere model. *J. Space Weather Space Clim.* **5**, A1. <https://doi.org/10.1051/swsc/2015001> (2015).
- Sutton, E. K., Cable, S. B., Lin, C. S., Qian, L. & Weimer, D. R. Thermospheric basis functions for improved dynamic calibration of semi-empirical models. *Space Weather* <https://doi.org/10.1029/2012SW000827> (2012).
- Doornbos, E., Klinkrad, H. & Visser, P. Use of two-line element data for thermosphere neutral density model calibration. *Adv. Space Res.* **41**, 1115–1122. <https://doi.org/10.1016/j.asr.2006.12.025> (2008).
- Storz, M. F., Bowman, B. R., Branson, M. J. I., Casali, S. J. & Tobiska, W. K. High accuracy satellite drag model (HASDM). *Adv. Space Res.* **36**, 2497–2505. <https://doi.org/10.1016/j.asr.2004.02.020> (2005).
- Boniface, C. & Bruinsma, S. Uncertainty quantification of the DTM2020 thermosphere model. *J. Space Weather Space Clim.* **11**, 53. <https://doi.org/10.1051/swsc/2021034> (2021).
- Licata, R. J., Mehta, P. M., Tobiska, W. K. & Huzurbazar, S. Machine-learned HASDM thermospheric mass density model with uncertainty quantification. *Space Weather* **20**, e2021SW002915. <https://doi.org/10.1029/2021SW002915> (2022).
- Marcos, F. et al. Precision low earth orbit determination using atmospheric density calibration. *J. Astronaut. Sci.* **46**, 395–409. <https://doi.org/10.1007/BF03546389> (1998).
- Nazarenko, A., Cefola, P. & Yurasov, V. Estimating atmospheric density variations to improve LEO orbit prediction accuracy. In *AIAA/AAS Space Flight Mechanics Meeting* (AAS 98-190, 1998). http://www.space-flight.org/AAS_meetings/1998_winter/abstracts/98-190.html.
- Bowman, B. & Storz, M. High accuracy satellite drag model (HASDM) review. In *AIAA/AAS Astrodynamics Specialist Conference* (AAS 03-625, 2003). https://sol.spaceenvironment.net/JB2008/pubs/JB2006_AAS_2003_625.pdf.
- Tobiska, W. K. et al. The SET HASDM density database. *Space Weather* <https://doi.org/10.1029/2020SW002682> (2021).
- Reigber, C., Lühr, H. & Schwintzer, P. Champ mission status. *Adv. Space Res.* **30**, 129–134 (2002).
- Mehta, P. M., Walker, A. C., Sutton, E. K. & Godinez, H. C. New density estimates derived using accelerometers on board the CHAMP and GRACE satellites. *Space Weather* **15**, 558–576. <https://doi.org/10.1002/2016SW001562> (2017).
- Mehta, P. M., McLaughlin, C. A. & Sutton, E. K. Drag coefficient modeling for grace using Direct Simulation Monte Carlo. *Adv. Space Res.* **52**, 2035–2051. <https://doi.org/10.1016/j.asr.2013.08.033> (2013).
- Mehta, P. M., Walker, A., McLaughlin, C. A. & Koller, J. Comparing physical drag coefficients computed using different gas-surface interaction models. *J. Spacecr. Rockets* **51**, 873–883. <https://doi.org/10.2514/1.A32566> (2014).
- Walker, A., Mehta, P. & Koller, J. Drag coefficient model using the Cercignani-Lampis-Lord gas-surface interaction model. *J. Spacecr. Rockets* **51**, 1544–1563. <https://doi.org/10.2514/1.A32677> (2014).
- Bettadpur, S. Gravity Recovery and Climate Experiment: Product Specification Document. *GRACE 327-720, CSR-GR-03-02* (2012). Cent. for Space Res., The Univ. of Texas, Austin, TX, <https://podaac.jpl.nasa.gov/GRACE>.

34. Drinkwater, M. R., Floberghagen, R., Haagmans, R., Muzi, D. & Popescu, A. *GOCE: ESA's First Earth Explorer Core Mission* 419–432 (Springer, 2003).
35. Friis-Christensen, E., Lühr, H. & Hulot, G. Swarm: A constellation to study the earth's magnetic field. *Earth Planets Space* **58**, 351–358. <https://doi.org/10.1186/BF03351933> (2006).
36. Yau, A. W. & James, H. G. CASSIOPE enhanced polar outflow probe (e-POP) mission overview. *Space Sci. Rev.* **189**, 3–14. <https://doi.org/10.1007/s11214-015-0135-1> (2015).
37. Calabia, A. & Jin, S. Upper-atmosphere mass density variations from CASSIOPE precise orbits. *Space Weather* **19**, e2020SW002645. <https://doi.org/10.1029/2020SW002645> (2021).
38. Calabia, A. & Jin, S. Thermospheric mass density disturbances due to magnetospheric forcing from 2014–2020 CASSIOPE precise orbits. *J. Geophys. Res. Space Phys.* **126**, e2021JA029540. <https://doi.org/10.1029/2021JA029540> (2021).
39. Weimer, D. R., Sutton, E. K., Mlynczak, M. G. & Hunt, L. A. Intercalibration of neutral density measurements for mapping the thermosphere. *J. Geophys. Res. Space Phys.* **121**, 5975–5990. <https://doi.org/10.1002/2016JA022691> (2016).
40. March, G., Doornbos, E. & Visser, P. High-fidelity geometry models for improving the consistency of CHAMP, GRACE, GOCE and Swarm thermospheric density data sets. *Adv. Space Res.* **63**, 213–238. <https://doi.org/10.1016/j.asr.2018.07.009> (2019).
41. Sheridan, P. L. *Updates and Improvements to the Satellite Drag Coefficient Response Surface Modeling Toolkit*. Master's thesis, West Virginia University, <https://doi.org/10.33915/etd.8227> (2021).
42. Ganushkina, N., Jaynes, A. & Liemohn, M. Space weather effects produced by the ring current particles. *Space Sci. Rev.* **212**, 1315–1344. <https://doi.org/10.1007/s11214-017-0412-2> (2017).
43. Weimer, D. R. *et al.* Improving neutral density predictions using exospheric temperatures calculated on geodesic, polyhedral grid. *Space Weather* **18**, e2019SW002355. <https://doi.org/10.1029/2019SW002355> (2020).
44. Licata, R. J., Mehta, P. M., Weimer, D. R. & Tobiska, W. K. Improved neutral density predictions through machine learning enabled exospheric temperature model. *Space Weather* **19**, e2021SW002918. <https://doi.org/10.1029/2021SW002918> (2021).
45. Iyemori, T. Storm-time magnetospheric currents inferred from mid-latitude geomagnetic field variations. *J. Geomagn. Geoelectr.* **42**, 1249–1265. <https://doi.org/10.5636/jgg.42.1249> (1990).
46. Siciliano, F. *et al.* Forecasting SYM-H index: A comparison between long short-term memory and convolutional neural networks. *Space Weather* **19**, e2020SW002589. <https://doi.org/10.1029/2020SW002589> (2021).
47. Matsuo, T. & Forbes, J. M. Principal modes of thermospheric density variability: Empirical orthogonal function analysis of CHAMP 2001–2008 data. *J. Geophys. Res. Space Phys.* <https://doi.org/10.1029/2009JA015109> (2010).
48. Lei, J., Matsuo, T., Dou, X., Sutton, E. & Luan, X. Annual and semiannual variations of thermospheric density: EOF analysis of CHAMP and GRACE data. *J. Geophys. Res. Space Phys.* <https://doi.org/10.1029/2011JA017324> (2012).
49. Mehta, P. M. & Linares, R. A methodology for reduced order modeling and calibration of the upper atmosphere. *Space Weather* **15**, 1270–1287. <https://doi.org/10.1002/2017SW001642> (2017).
50. Mehta, P. M., Linares, R. & Sutton, E. K. A quasi-physical dynamic reduced order model for thermospheric mass density via hermitian space-dynamic mode decomposition. *Space Weather* **16**, 569–588. <https://doi.org/10.1029/2018SW001840> (2018).
51. Gondelach, D. J. & Linares, R. Real-time thermospheric density estimation via two-line element data assimilation. *Space Weather* **18**, e2019SW002356. <https://doi.org/10.1029/2019SW002356> (2020).
52. Pearson, K. L. I. I. On lines and planes of closest fit to systems of points in space. *Lond. Edinb. Dublin Philos. Mag. J. Sci.* **2**, 559–572. <https://doi.org/10.1080/14786440109462720> (1901).
53. Hotelling, H. Analysis of a complex of statistical variables into principal components. *J. Educ. Psychol.* **24**, 417–441. <https://doi.org/10.1037/h0071325> (1933).
54. Licata, R. J., Mehta, P. M., Tobiska, W. K., Bowman, B. R. & Pilinski, M. D. Qualitative and quantitative assessment of the SET HASDM database. *Space Weather* **19**, e2021SW002798. <https://doi.org/10.1029/2021SW002798> (2021).
55. Wang, S., Weng, L., Fang, H., Xie, Y. & Yang, S. Intra-annual variations of the thermospheric density at 400 km altitude from 1996 to 2006. *Adv. Space Res.* **54**, 327–332. <https://doi.org/10.1016/j.asr.2013.12.011> (2014).
56. Weng, L., Lei, J., Zhong, J., Dou, X. & Fang, H. A machine-learning approach to derive long-term trends of thermospheric density. *Geophys. Res. Lett.* **47**, e2020GL087140. <https://doi.org/10.1029/2020GL087140> (2020).
57. Kodikara, T., Carter, B. & Zhang, K. The first comparison between swarm-c accelerometer-derived thermospheric densities and physical and empirical model estimates. *J. Geophys. Res. Space Phys.* **123**, 5068–5086. <https://doi.org/10.1029/2017JA025118> (2018).
58. Chen, H., Liu, H. & Hanada, T. Storm-time atmospheric density modeling using neural networks and its application in orbit propagation. *Adv. Space Res.* **53**, 558–567. <https://doi.org/10.1016/j.asr.2013.11.052> (2014).
59. Choury, A., Bruinsma, S. & Schaeffer, P. Neural networks to predict exosphere temperature corrections. *Space Weather* **11**, 592–602. <https://doi.org/10.1002/2013SW000969> (2013).
60. Linnainmaa, S. *The representation of the cumulative rounding error of an algorithm as a Taylor expansion of the local rounding errors*. Master's thesis, University of Helsinki (1970).
61. Dreyfus, S. The computational solution of optimal control problems with time lag. *IEEE Trans. Autom. Control* **18**, 383–385. <https://doi.org/10.1109/TAC.1973.1100330> (1973).
62. O'Malley, T. *et al.* Keras Tuner. <https://github.com/keras-team/keras-tuner> (2019).
63. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I. & Salakhutdinov, R. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **15**, 1929–1958 (2014).
64. Alake, R. Understanding and Implementing Dropout in TensorFlow and Keras. <https://towardsdatascience.com/understanding-and-implementing-dropout-in-tensorflow-and-keras-a8a3a02c1bfa> (2020).
65. Gal, Y. & Ghahramani, Z. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning (2016). arXiv: [1506.02142v6](https://arxiv.org/abs/1506.02142v6).
66. Nix, D. & Weigend, A. Estimating the mean and variance of the target probability distribution. In *Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN'94)*, vol. 1, 55–60. <https://doi.org/10.1109/ICNN.1994.374138> (1994).
67. Licata, R. J., Tobiska, W. K. & Mehta, P. M. Benchmarking forecasting models for space weather drivers. *Space Weather* **18**, e2020SW002496. <https://doi.org/10.1029/2020SW002496> (2020).

Acknowledgements

P.M.M. gratefully acknowledges support under NSF CAREER award #2140204. The authors would like to acknowledge DLR for their work on the CHAMP mission along with GFZ Potsdam for managing the data.

Author contributions

P.M.M. created the CHAMP dataset, and R.J.L. prepared the HASDM PCA dataset. R.J.L. and P.M.M. developed models and analyzed the results. All authors reviewed the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-022-11049-3>.

Correspondence and requests for materials should be addressed to R.J.L.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022