

Random forest model of ultra-low frequency magnetospheric wave power

Article

Published Version

Creative Commons: Attribution 4.0 (CC-BY)

Open Access

Bentley, S. N., Stout, J., Bloch, T. E. and Watt, C. E. J. (2020) Random forest model of ultra-low frequency magnetospheric wave power. *Earth and Space Science*, 7 (10). e2020EA001274. ISSN 2333-5084 doi: <https://doi.org/10.1029/2020EA001274> Available at <https://centaur.reading.ac.uk/92475/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1029/2020EA001274>

Publisher: American Geophysical Union

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

Earth and Space Science



RESEARCH ARTICLE

10.1029/2020EA001274

Random Forest Model of Ultralow-Frequency Magnetospheric Wave Power

S. N. Bentley¹ , J. R. Stout¹ , T. E. Bloch¹ , and C. E. J. Watt^{1,2}

¹Department of Meteorology, University of Reading, Reading, UK, ²Department of Mathematics, Physics and Electrical Engineering, Northumbria University, Newcastle upon Tyne, UK

Key Points:

- Decision tree ensembles of ULF wave power are presented; these solar wind-driven higher-resolution models outperform previous models
- Models are used to test hypotheses of magnetic local time wave occurrence, such as dawn-dusk asymmetries and conditions that vary with B_z
- Remaining uncertainty in the model suggests that the internal state of the magnetosphere must be characterized in next-generation models

Correspondence to:

S. N. Bentley,
snbentley@outlook.com

Citation:

Bentley, S. N., Stout, J. R., Bloch, T. E., & Watt, C. E. J. (2020). Random forest model of ultralow-frequency magnetospheric wave power. *Earth and Space Science*, 7, e2020EA001274. <https://doi.org/10.1029/2020EA001274>

Received 26 MAY 2020

Accepted 10 AUG 2020

Accepted article online 15 AUG 2020

Abstract Models of magnetospheric ultralow-frequency (ULF) waves can be used to study wave phenomena and to calculate the effect of these waves on the energization and transport of radiation belt electrons. We present a decision tree ensemble (i.e., a random forest) model of ground-based ULF wave power spectral density driven by solar wind speed v_{sw} , north-south component of the interplanetary magnetic field B_z , and variance of proton number density $\text{var}(N_p)$. This model corresponds to four radial locations in the magnetosphere (roughly $L \sim 4.21$ to 7.94) and spans 1–15 mHz, with hourly magnetic local time resolution. The decision tree ensembles are easier to use than the previous model generation; they have better coverage, perform better at predicting power, and have reduced error due to intelligently chosen bins in parameter space. We outline the difficulties in extracting physics from parameterized models and demonstrate a hypothesis testing framework to iteratively explore finer driving processes. We confirm a regime change for ULF driving about $B_z = 0$. We posit that ULF wave power directly driven by magnetopause perturbations corresponds to a latitude-dependent dawn-dusk asymmetry, which we see with increasing speed. Model uncertainty identifies where the underlying physics is not fully captured; we find that power due to substorms is less well characterized by $B_z > 0$, with an effect that is seen globally and not just near midnight. The largest uncertainty is seen for the smallest amounts of solar wind driving, suggesting that internal magnetospheric properties may play a significant role in ULF wave occurrence.

Plain Language Summary To describe the energy and motion of electrons in the radiation belts for space weather modeling purposes, we need to be able to describe the large-scale ultralow-frequency plasma waves. These waves can energize and transport high-energy electrons, which pose a hazard to spacecraft. Using ground-based magnetometers whose observations directly correspond to waves in space, we construct a machine learning model that predicts power in those waves based on incoming solar wind properties. This performs better than the previous model and is more easily used to investigate the physics. We demonstrate how to use it to explore the physics by comparing it to existing wave phenomena. We find that the solar wind drives it in the ways expected but that we should not average over these processes as some of them have a significant on/off switch. We find that two aspects of magnetospheric physics are not fully captured in our model and contribute to the remaining uncertainty: the effect of substorms throughout the magnetosphere and the internal state of the magnetosphere. These are likely to have a significant additional impact on wave occurrence.

1. Introduction

Ultralow-frequency (ULF) plasma waves are implicated in the energization, transport, and radial diffusion of electrons in Earth's radiation belts (Elkington et al., 1999; Ma et al., 2018; Mann et al., 2012). As such, global characterization of these waves is necessary for global radiation belt models, in order to nowcast and forecast an environment that is hostile to the satellites underpinning our everyday life. In particular, diffusion-based radiation belt models such as the British Antarctic Survey Radiation Belt Model (BAS-RBM; Glauert et al., 2014), Versatile Electron Radiation Belt (VERB; Subbotin et al., 2010), and the Dynamic Radiation Environment Assimilation Model (DREAM; Reeves et al., 2012) utilize a radial diffusion coefficient that depends on the power in magnetospheric ULF waves. Horne et al. (2013) identified improved characterization of ULF waves throughout the magnetosphere as one of the next steps necessary for such diffusion-based models.

©2020. The Authors.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

In addition to radiation belt forecasting, ULF wave models can be used to examine the underlying physics of wave generation and propagation. The parameters chosen for either type of model should be causally related to ULF wave power spectral density (PSD). Moreover, forecasting models should only use those parameters that are routinely available. Bentley et al. (2018) found that to satisfy these criteria, one must account for interdependence between solar wind properties that can result in misleading correlations. They concluded that the three dominant solar wind parameters driving magnetospheric ULF waves observed at the ground were the speed v_{sw} , interplanetary north-south magnetic field component B_z , and the summed power in the proton number density δNp (which can be more simply represented as the variance $\text{var}(Np)$). Bentley et al. (2019) used these parameters to construct a simple empirical model of ULF wave power at several locations, including an azimuthal (magnetic local time, MLT) resolution to account for the variation in power around the Earth. (For brevity, we will refer to PSD as simply “power” throughout this paper.)

The model proposed by Bentley et al. (2019) is a first attempt at including the uncertainty or variance in a ULF model, just as uncertainty is included successfully in many weather and climate modeling techniques (see, e.g., Berner et al., 2017; Langan et al., 2014; Leutbecher et al., 2017; Weisheimer et al., 2011). However, in order to investigate uncertainty, the Bentley et al. (2019) model divided the input parameter axes into linearly spaced bins. Large linearly spaced bins are not optimal for data analysis or time series prediction as the resulting model will be susceptible to several sources of bias. In this paper, we use a well-known machine learning method, an ensemble of regression decision trees (i.e., a random forest), to construct a model of ULF wave PSD that does not use large linearly spaced bins. Decision trees allow for a data-driven approach to binning the feature space of input data. The resulting variable size bins (compared to Bentley et al., 2019) give us better resolution for changes in ULF power with respect to the underlying parameters. Hence, the model presented here has better coverage, more accurate predictions, reduced bias, and the ability to include more parameters with the same amount of data. This allows us to improve the MLT resolution, providing a more useful prediction and the ability to examine the effect of each solar wind parameter on ULF wave power with MLT.

With such statistical or parameterized models, the ability to investigate the underlying physics varies not only with the choice of parameters but also the interdependence of those parameters and the physical processes. This is because parameterized models approximate a predicted output value such as ULF wave power as a hypersurface, with all physical processes compounded together. In this situation, assuming that the effect of each parameter can be used to unambiguously trace the underlying physical processes is equivalent to assuming that the effect of each process adds linearly to the final output (here, ULF wave power). When the input parameters are interdependent, and the physical processes are not only dependent on each other but also upon multiple input variables, the resulting model can be highly nonlinear and a more circumspect approach must be taken. Here we suggest testing a series of hypotheses addressing successively finer processes rather than attempting to tease out all physical processes at once. We show how to construct the starter hypotheses for the simple case of variable power with MLT and suggest further processes affecting each one that may be investigated in future.

We begin by explaining the data processing and the theory behind decision trees in section 2. In section 2.1 we describe and construct the decision trees. In section 2.2 we test the decision tree ensembles and compare their predictions to the previous model in Bentley et al. (2019), including discussing how it reduces the error seen in the initial model. In section 3 we show an example of how such models can be used to test our understanding of the underlying physics; we hypothesize where we would expect to see wave power change with MLT based on existing theories and then compare this to both the variation of power and the location of the remaining uncertainty. We discuss how the use of decision tree ensembles mitigates known issues with parameterizations in space physics in section 4 and conclude in section 5.

2. Constructing the Ensemble of Decision Trees (the Random Forest)

Our model is constructed using 15 years (1990–2005) of solar wind observations from OMNIWeb (King & Papitashvili, 2005) and ground-based magnetometer observations from the CANOPUS magnetometer chain, now upgraded to CARISMA (Mann et al., 2008; Rostoker et al., 1995). From OMNIWeb, we use hourly values of v_{sw} and B_z , while $\text{var}(Np)$ is constructed from 1 min data, interpolating data gaps if 8 min or fewer are missing per hour and discarding that hour otherwise. From the magnetometer chain, we use ground stations FCHU, GILL, ISLL, and PINA, which have average L values of 7.94, 6.51, 5.40, and 4.21 over the 15 years

used (values from Rae et al., 2012). The processing method is the same as used in Bentley et al. (2019). To summarize, processing involves removing outliers from the 5 s resolution data and then calculating PSD in hourly windows. After using first a low-pass Butterworth filter with passband 0–84 mHz, the multitaper method with a time halfbandwidth product of 1.4 is applied to find the PSD at each frequency. Details and justification of the PSD calculation can be found in Bentley (2019). Ensembles of decision trees are constructed (and made available) for each of these frequencies, but we choose to show all our examples at 5 mHz throughout this paper.

Note that throughout this work we refer to solar wind properties and other inputs as *input parameters* and to the model settings as either *hyperparameters* or simply *settings*. In machine learning terminology these would usually be called *input features* and (*hyper*)*parameters*, respectively, while our predicted PSD output would be called the *target feature*.

For each station and frequency, we train a model on the input parameters $[v_{sw}, Bz, \text{var}(Np), MLT]$. The three solar wind parameters were identified by Bentley et al. (2018) as the three dominant drivers of instantaneous ULF wave power. We convert MLT from a linear axis $m = 0\text{--}23$ to a cyclic one, that is, $m_x = \sin\left(\frac{2\pi m}{24}\right)$, $m_y = \cos\left(\frac{2\pi m}{24}\right)$. This removes an artificial split at midnight, allowing for continuous MLT. We train a model at each station, frequency, and component in order to retain the original discrete values as it is not appropriate to represent power continuously along these input parameters. Stations are by design discrete and spaced such that although power between them may be correlated, they should not be experiencing the effect of the same waves. Geomagnetic components represent an orthogonal basis and are usually considered to represent independent processes. This is idealized, but we expect this is a feature necessary for most users. Finally, frequency values should be separated as power can be smeared across nearby frequencies during the calculation of PSD. To account for this, users should either integrate total power of neighboring frequencies or select frequencies with the appropriate spacing. This effect is removed here by showing all results at a single frequency (5 mHz) and is discussed in the documentation of the saved models for any interested users.

We choose a decision tree as we can not only improve the predictions of Bentley et al. (2019) but also drastically improve the resolution of bins in the parameter space, allowing us to examine changes of power with MLT. The previous model used causally correlated parameters to reduce the bias resulting from the interdependent relationships of solar wind properties, but other sources of bias and variance remained, such as the choice of bins. Using large linearly spaced bins results in a compromise between defining large enough bins in those parts of parameter space that are sparsely sampled and having too large a bin full of data points that behave quite differently from one end of the bin to another. Space physics observations are often not uniformly spread throughout the parameter space; for example, Bz observations cluster around $Bz = 0$. Applying large linear bins to Bz observations will result in bins that are better populated at one end, hence skewing any predictions to that edge of the bin. On the other hand, decreasing bin size would quickly result in bins with too small populations and hence poor coverage of the parameter space. Both these issues would result in skewing predicted wave power values to areas of the solar wind input parameter space ($[v_{sw}, Bz, \text{var}(Np)]$) where we have the most data. In this work, decision trees (and random forests) are applied to directly combat these limitations using variable bins determined by the variance of the data.

We will show that such models reduce the error compared to the previous model (using forecasting skill), yet are so much more efficient at using data that we gain drastically better resolution in MLT, using 24 MLT bins rather than just four sectors. All decision tree ensembles can be found on Zenodo, along with example code to read in and apply them to solar wind values (Bentley et al., 2020). The saved models include both geomagnetic east-west and north-south components. In the remainder of this paper we show results exclusively using the geomagnetic north-south models.

2.1. Determining the Decision Tree Ensemble Constraints

A basic decision tree is a relatively simple algorithm (Hastie et al., 2009). For regression problems predicting a continuous output from a continuous input it runs across the input parameter space (or “feature space”) to find where to make a split (i.e., a bin edge). Splits are chosen based on reducing the mean square error (MSE). Examples are the following:

1. The algorithm tests each input feature for partition locations, choosing the input parameter and parameter value that most reduces MSE of ULF PSD in the two newly partitioned regions (e.g., $\sim 450 \text{ km s}^{-1}$)




Figure 1. Visualizations of decision trees. (a) A representation of one of our decision trees, where each split becomes a branch. (b) An example of a regression decision tree, where the model in black approximates a continuous function fitted to the data in blue. The underlying function is $y = x^2$ plus a normally distributed error η . (c) A two-dimensional visualization of one of the trees in this paper, with bins overlying the original input data points.

2. After this “branch” our model has two bins, split by a threshold of $v_{sw} \sim 450 \text{ km s}^{-1}$. In each bin, the predicted ULF PSD will be approximated by the mean value of power observed in that bin.
3. This process is repeated until the model reaches specified hyperparameters such as the branch depth or the minimum allowed number of bins.

Some simple examples are shown in Figure 1; in (a) we see a dendrogram, a pictorial representation of one of the decision trees here. In (b) we see an example decision tree regression on one-dimensional data, made from a curve $y = x^2$ plus a normally distributed error η . This essentially fits a stepwise function with varying step sizes. In (c) we see an example decision tree such as the ones trained throughout this paper to predict ULF wave PSD but displayed only along solar wind speed v_{sw} and magnetic field component B_z . The bin size is clearly very variable, with many smaller bins capturing the rapid change of behavior around $B_z = 0$. In (c) lower power is indicated by dark blue and high power by light green. The color bar has been omitted to prevent misinterpretations from many busy overlays; analysis of power throughout the model can be found in sections below.

Decision trees tend to be low bias (as they make few assumptions about the model) but high variance (as they are strongly affected by variance in the input data set) (Mehta et al., 2019). Decision trees are particularly prone to overfitting, making so many small bins that they do not generalize well to nontraining data. We can reduce these effects using hyperparameter optimization and ensembles of decision trees to make a model that does generalize well.

Hyperparameters are settings that define the model architecture by constraining the final tree, for example, by determining a set number of bins (leaves) the tree may end up with. Figure 2a is a cartoon illustrating the importance of using the correct settings in a general machine learning model. This is shown for a theoretical hyperparameter (or degree of freedom in the model). An example hyperparameter for another type of model would be what order polynomial to use when curve fitting. As this theoretical hyperparameter increases from 0, the MSE in the model when applied to test data drastically drops but then increases again as overfitting occurs. Similarly, we must perform optimization to choose the correct degrees of freedom for our model: We choose to specify the minimum samples per leaf (i.e., in each final bin) and the maximum depth of each tree. To choose these, we trained multiple models using permutations of these settings, calculating the MSE using maximum depths between 5 and 35, and minimum samples in a leaf between 5 and 100. (Input values were chosen to appropriately and efficiently cover the hyperparameter space. While we sampled every maximum depth value between 5 and 35, we calculated MSE for every minimum samples-per-leaf value between 5 and 10, every second value between 10 and 20, every five values up to 50, and then every tenth value up to 100.) To reduce overfitting in the final model from our hyperparameter choice, fivefold cross validation can be used to calculate the MSE estimate for each pair of hyperparameters tested. Using this method, the training data are partitioned into five segments. Then a model is trained on four segments, and the MSE is calculated on the withheld segment. This process is iterated across the segments until each segment has been both part of the training and test sets, so that the final MSE values for each pair of hyperparameter settings are the average of five total runs. These MSE values were used to rank the possible settings, as shown in Figure 2b. We concluded that the optimum model has a maximum depth of 11 and a minimum number of samples of 18 per leaf.




Figure 2. Choosing settings (hyperparameters) for a machine learning model. (a) A cartoon example of any hyperparameter for any machine learning model; a simple example is the order of a polynomial used in curve fitting. As this example parameter increases, mean square error on the test data reduces as the model improves, but then rises again as overfitting occurs. The ideal value for this parameter would then be the minimum point. (b) The results of hyperparameter optimization for our decision trees. Decision trees with the given minimum samples per leaf and maximum depth were ranked to indicate those with the lowest mean square error.

Finally, while we talk about “decision tree” models throughout this paper, each final model is really an averaged decision tree created from an ensemble, known as a random forest (Hastie et al., 2009). The final output is the average (mean) value output from all ensemble members, and the members vary by training each on different collections of the training data. Each collection is the same size as the original training data, where samples are drawn with replacement. We use 256 trees in each ensemble, a value chosen to be computationally manageable yet retain enough members that the forest should be stable with the addition of new trees. (Hastie et al., 2009; Oshiro et al., 2012). To find this averaged value for a given input, we take the 256 predicted values output by each tree and take the mean. Thus, the final model—the random forest—does not easily decompose into bins, yet it is constructed from a set of models that do, individually, have the desired properties in parameter space via variable bins. Averaging over this ensemble of decision trees results in a final model that has reduced variance and therefore will generalize far better to new data.

2.2. Testing the Decision Tree Ensembles

To construct a general machine learning model, the data set should be split into testing and training data. The larger the training data set, the easier it is to make a model with reduced variance. However, this would mean less data to test on and therefore a performance statistic with more variance.

During the hyperparameter optimization, this train:test split was used as part of the fivefold cross validation. We use a similar method to test our model against overfitting. For each station, frequency, and component, we train 100 models using 80% of the data and calculate the MSE on the remaining test data, taking the median MSE across the hundred runs. We find that over all stations and frequencies, the smallest and largest median MSE values on the test partitions were 0.13 and 0.68 $\log_{10}(\text{PSD})$, $(nT)^2/\text{Hz}$, respectively.

From this point onward we use decision tree ensembles retrained on the full data set for science and forecasting and in order to compare to the previous model. We quantify the improvement in predictions using forecasting skill,

$$\text{Skill} = 100 \left(1 - \frac{\text{MSE}_1}{\text{MSE}_2} \right), \quad (1)$$

which compares the MSE between a model and observations to a reference model and observations, so that positive values indicate the tested model is better than the reference model. As in Bentley et al. (2019), we use a random model as the reference. In Table 1 we reproduce the forecasting skill for selected frequencies and stations, including the decision tree model.

The decision tree ensembles consistently outperform other models at each station and frequency. While this increase in skill is relatively small, the decision tree models are also far more suitable for analysis given their efficient coverage of parameter space. This means that we can use model-represented power to approximate

Table 1
Forecasting Skill at Selected Stations and Frequencies

Station/frequency tested	Model skill score versus random reference model			
	24 hr persistence	1 hr persistence	Previous model (means only)	Decision tree
FCHU/3.06 mHz	35.1	69.2	74.6	78.2
GILL/3.33 mHz	37.5	73.9	78.0	81.2
ISLL/4.17 mHz	37.4	76.1	78.4	81.6
PINA/4.44 mHz	35.5	72.7	77.6	80.4

Note. Forecasting skill scores for four stations and frequencies, testing the ability of each parameterized model to reproduce the original 15 years of data. The baseline reference model used is a “random” model, where power is sampled from the original total distribution of the given partition. Simple 24 and 1 hr “persistence” models are tested against this baseline (i.e., assuming power in the oncoming hour is the same as the previous day or hour) in addition to the solar wind-parameterized model. The probability distributions predicted for each hour by the solar wind model were either sampled or the mean value was taken to construct each 15 year time series. Where sampling methods were used, 2,000 time series were made and the forecast skill calculated for each one; the median is shown here. (Some skill scores vary here from those shown in Bentley et al., 2019. These small [~ 0.1] variations are in 24 and 1 hr persistence scores and are due to the random model used in skill score calculations. Both the previous model and decision tree scores converge well before 2000 runs; lower skill scores appear more susceptible to the random aspect of this calculation.)

power throughout the magnetosphere and test hypotheses about when and where we see wave power, particularly with changing MLT. We can also examine the remaining uncertainty in the model; this is included as part of the MLT analysis in section 3.

While we do not consider how to use the decision trees probabilistically in this paper, we can still test how well it reproduces the entire distribution over the 15 years. We show this for four MLTs, shown in Figure 3. Each plot shows the overall quantiles of power seen at GILL 5 mHz over the 15 year period, both the original observed values and the reproductions using our decision tree ensembles. The original values are shown in brown with crosses and the reproductions with blue circles. In all sectors, the median part of the distribution is captured well, but the model underrepresents the extreme high and low powered values. Since the decision tree ensembles are trained using MSE, this is to be expected. We may be able to resolve this in future by using the underlying distributions in each bin. However, for the remainder of this paper we focus on showing how these decision tree models can be used to test existing hypotheses on ULF wave generation and propagation and their variation by MLT and solar wind conditions.

3. Using the Models to Test Hypotheses: Wave Power Variation With MLT

There already exist many statistical studies of varying ULF power with MLT that identify asymmetries in wave power and occurrence (see, e.g., Baker et al., 2003; Berube et al., 2014; Pahud et al., 2009). In most such statistical studies, investigators must choose driving parameters to focus on and therefore manually restrict other parameters. In contrast, our model chooses the appropriate resolution for each and so we have more dimensions on which to investigate the underlying physics. Given that our characterization of power has six dimensions (when one includes latitude and frequency) and is likely to contain abundant physics, we do not attempt to discuss all the constituent processes here. Indeed, since it is unlikely that we will be able to extract every single process contributing to wave power, we begin with simple testable hypotheses to determine the dominant driving processes and suggest that to investigate the remaining processes and power, successively finer hypothesis be constructed. This framework of successive questions should (a) state the process we think will dominate observations, (b) propose how this process would manifest in the model, (c) test for observations confirming or contradicting these expectations, and (d) repeat for finer processes. We draw on previous literature to make hypotheses about the occurrence of ULF wave power in our models and how this varies with MLT. We then test these hypotheses by examining power in our model. Despite choosing to only test MLT-dependent hypotheses, we find that to examine these, we must involve some power variations across other input parameters. Nevertheless, extending hypothesis testing to fully include power dependence on other parameters is out of the scope of this model introduction paper. This section remains an example of how to use the model we have presented and quantified throughout the paper. Thus, we do not examine power variation with frequency as none of the MLT hypotheses we test require this. Instead, we choose one representative frequency (5 mHz) for this section and simply note that similar signatures can be seen at other

frequencies. A full study of the variation of power with parameters not the focus of this section—including frequency—is left for future work.

3.1. Hypotheses: Variations of ULF Wave Power

We focus on testable hypotheses of power asymmetries in MLT, drawing from results of numerous previous studies (e.g., Pahud et al., 2009; Rae, 2017; Takahashi et al., 2016). Many of these previous studies were based on in situ (rather than ground) observations or investigated ULF wave events such as field line resonances. In contrast, our model does not distinguish between power contributions from broadband or narrowband wave activity and does not use any criteria or thresholds to define ULF wave events. This may result in different patterns of MLT asymmetry for ground broadband power than for, say, in situ field line resonances. We will investigate whether our existing understanding holds in our model and whether MLT asymmetries can tell us what driving processes our parameters v_{sw} , B_z , and $\text{var}(Np)$ represent.

Hypothesis 1: There will be different asymmetries in wave power with MLT for $B_z < 0$ compared to $B_z > 0$.

Testing this hypothesis will inform us whether there is need to separate $B_z \pm 0$ throughout the analysis. Many previous studies have identified B_z as a significant driving parameter (see, e.g., Mathie & Mann, 2001; Simms et al., 2010; Takahashi et al., 2012), although often a simple linear relationship between ULF wave power and B_z was assumed. Bentley et al. (2018) found that ULF wave power dependence on B_z changed across the threshold $B_z = 0$.

Hypothesis 2: Waves driven by external magnetopause perturbations, particularly large amplitude ones, will have more power in the dawn sector. This asymmetry will be larger at higher-latitude stations.

The dawn asymmetry of ULF wave occurrence has been identified by many studies (see, e.g., Baker et al., 2003; Takahashi et al., 2016; Wright et al., 2018). This is thought to be due to the superior ability of the dawn side to support externally driven ULF waves, as there exists an underlying plasma mass density asymmetry. The radial gradient of plasma mass density in dawn is shallower than on the dusk side, where there are often also density structures such as plumes. On average this means that standing waves are less likely to form on the duskside (see, e.g., Rae, 2017, and references therein). This contribution is one of the easiest to test using our model and is therefore the option we will study first. Alternatively, the dawn-dusk asymmetry may be due to conditions that would roughly equally affect all external drivers, such as asymmetric magnetopause instabilities (Hwang & Sibeck, 2016; Mann & Wright, 1999) or the angle of the solar wind striking the dawn flank. Ideally in future we would also quantify the relative contribution of each of these conditions (likely




Figure 3. Power spectral density quantiles for observations (brown crosses) and reproductions using the decision tree model (blue circles) for GILL station at 5 mHz over the entire 15 years of input data. Quantiles are shown for a single hour in each of the midnight, dawn, noon, and dusk sectors.

requiring multiple further hypotheses), and for the moment, we note that it is believed that solar wind striking angle may have a lesser effect (Rae, 2017) and the effect of solar wind orientation was not significant enough to be resolved when choosing causal parameters in Bentley et al. (2018).

To test the radial density gradient explanation, we investigate whether power asymmetries manifest with the expected signature in our predictive model. We expect any dawn-dusk asymmetry to be larger for higher-latitude stations. Externally driven waves cannot penetrate as far into the magnetosphere, and therefore this asymmetry will drop off with decreasing station latitude. Rae et al. (2019) show that this “penetration distance” varies significantly in storms as the density changes, and therefore we expect that there will be a gradual lessening of this asymmetry with decreasing latitude, to reflect the variable nature of the plasma density.

However, we cannot fully test the signature of ULF wave power driven by magnetopause perturbations without isolating those perturbations, which we cannot do with our data set. Therefore, to test this hypothesis, we apply it separately to each of the solar wind parameters. If we see the dawn-dusk, latitudinal asymmetry signature for some individual parameters that are considered to be a proxy for magnetopause oscillations, we can conclude that the signature fulfills Hypothesis 2, suggesting that different radial density profiles is the dominant distinguishable cause of a dawn-dusk asymmetry in ULF wave power. On the other hand, if we see this asymmetry signature for all parameters, then we will not have proven Hypothesis 2 as there are multiple explanations to such an outcome. In that case, either *all* parameters correspond to driving via magnetopause perturbations, or the dawn-dusk asymmetry does not relate to the ability of magnetopause-driven waves to penetrate the magnetosphere at all but instead reflects other global asymmetries such as magnetopause instabilities, the solar wind striking angle, or the transmission of ULF waves through the ionosphere (which was omitted during the construction of this hypothesis and is discussed in the analysis, section 3.2.2).

For each parameter we also specify alternative hypotheses, to distinguish contributions from driving processes that do not involve magnetopause perturbations.

Hypothesis 2a: If speed is a proxy for wave power driven by Kelvin-Helmholtz instabilities and their resultant magnetopause perturbations, then power will increase more at dawn than dusk, particularly at high-latitude stations.

Alternatively, power everywhere increases with increasing speed.

We particularly expect to see this asymmetry with solar wind speed v_{sw} as this is often used as a proxy for the amount of driving by Kelvin-Helmholtz instabilities on the magnetopause. In this case, increasing speed would be associated with an increasing amount of power in the dawn sector, particularly at higher-latitude stations. An alternative finding may be that total power increases everywhere with speed in the same manner, suggesting that speed is a good proxy for the total energy entering the magnetosphere.

Hypothesis 2b: If B_z represents magnetopause perturbations such as those from traveling flux tubes, power will increase in dawn (particularly at high latitudes) for more strongly negative $B_z < 0$.

Alternatively, power may increase in the midnight sector with more negative B_z due to substorm activity.

In line with previous literature (see Bentley et al., 2018, and references therein), the increase of wave power for $B_z < 0$ (i.e., below the threshold $B_z = 0$) may be due to the Dungey cycle of magnetospheric reconnection: Power increases either due to perturbations along the magnetopause due to traveling flux tubes and/or due to substorms. Traveling flux tubes would result in the dawn-dusk asymmetries discussed above, whereas substorm-related increases in power would be expected to peak in the midnight sector. Note that there are multiple other processes related to B_z that impact ULF wave power, such as the magnetopause location (discussed below) and the properties of the low-latitude boundary layer (affecting the Kelvin-Helmholtz stability of the magnetopause and the radial Alfvén speed profile, and therefore ULF wave propagation) (Hwang & Sibeck, 2016). It is unlikely that we will be able to distinguish the impact of these from our three solar wind parameter model. Therefore, we choose to test hypotheses based on what previous work suggests may be the most dominant processes but will consider the possible impact of alternatives when analyzing the results.

Hypothesis 2c: If $\text{var}(N_p)$ represents solar wind buffeting of the magnetopause directly driving compressional magnetospheric waves, we will see the signature of higher power in dawn and at high latitudes.

Alternatively, $\text{var}(Np)$ represents unknown, complex impacts of sudden global compressions of the magnetopause.

Variations in either density or pressure have been associated with increased ULF wave power in multiple studies (see, e.g., Berube et al., 2014; Simms et al., 2010; Takahashi et al., 2012). Solar wind pressure or density variations have been shown to drive magnetospheric compressional waves and set up standing waves (Claudepierre et al., 2010; Wright & Rickard, 1995), and it has been shown that power is controlled by variations of pressure rather than by speed at latitudes lower than $L \sim 6$ (Berube et al., 2014; Takahashi et al., 2012). Clearly, pressure and/or number density have some control over the ULF wave power observed throughout the magnetosphere. However, it is unclear how much of this control is attributable to (1) the inherent properties of a globally compressed magnetosphere, (2) to many smaller plasma variations buffeting the magnetosphere, or (3) to the effects of sudden global compressions. (For the purposes of discussion we are decoupling the state of a compressed magnetosphere from the sudden compression itself as this helps us address processes individually; a constantly compressed magnetosphere may affect ULF waves differently to a suddenly compressed one, even though these states are related.) We can discuss these three possibilities in order.

First, in a compressed magnetosphere ULF waves can consistently propagate further, reaching lower-latitude stations due to the reduced distance to the magnetopause. Second, if $\text{var}(Np)$ represents the onset of particularly variable plasma, it may be that this buffeting drives magnetopause perturbations and hence radially propagating compressional waves. These may be due to the magnetosheath amplifying variations in number density or to magnetosheath processing driving other foreshock phenomena (see, e.g., Archer et al., 2013; Hartinger et al., 2013; Zastenker et al., 2002). Third, the $\text{var}(Np)$ contribution to ULF wave power may be due to the effects of sudden global compressions. This would include impacts such as the global scale “ringing” of the magnetosphere or the sudden loss of magnetospheric plasma to the solar wind, which would drastically change the radial profile of the underlying plasma such that density at dawn and dusk are more symmetric.

Based on the three solar wind parameters in our model, we would expect interplanetary shocks to be primarily represented by $\text{var}(Np)$. Indeed, Bentley et al. (2018) concluded that $\text{var}(Np)$ represented large solar wind structures sweeping across Earth's magnetosphere (i.e., shocks) rather than solar wind MHD waves but were unable to speculate on the intermediate processes driving higher power in magnetospheric waves. By breaking down the possible contributions of $\text{var}(Np)$ as we have in the previous paragraph (i.e., via a static compressed state/variable solar wind plasma driving magnetopause perturbations/ sudden global compressions), we can more easily distinguish the underlying processes affecting ULF waves, as shocks could be affecting ULF wave power through any combination of the above processes, all of which would have varied impacts across dayside wave power.

Ideally, we would be able to investigate the impact of all of these processes. However, due to the complexity of these processes, we cannot easily suggest how their combined effect would manifest in our model. This is where our method of successive hypothesis testing comes in to play: We will first examine the easiest process to detect. Therefore, our hypothesis is that $\text{var}(Np)$ represents some type of magnetopause perturbation (whether this is direct buffeting or involves magnetosheath processing), and our alternative hypothesis is that $\text{var}(Np)$ does not represent power by magnetopause perturbations but by some other complex interaction of the above processes during a sudden global compression. If this is the hypothesis accepted in our analysis, then we expect that future work will be needed to determine whether any of the remaining underlying processes can instead be shown to dominate $\text{var}(Np)$ contributions to power.

The missing, untestable hypothesis: When the magnetopause is compressed (i.e., for high solar wind speed, strongly negative B_z and large $\text{var}(Np)$), we will observe greater power on the dayside.

The equilibrium position of the magnetopause has been shown to have a clear impact on ULF wave power, with more power for a magnetopause nearer to the Earth (Murphy et al., 2015). The position of the magnetopause is strongly related to solar wind v_{sw} and B_z (Shue et al., 1998) and is likely to be indirectly related to $\text{var}(Np)$, which scales with Np . However, this hypothesis is very difficult to test with our model, because we can only capture the magnetopause location indirectly. Therefore, any contribution from magnetopause distance would be inextricably tied up with all the other interparameter relationships and with the processes they are already known to represent. For example, we already expect high solar wind speed to drive high

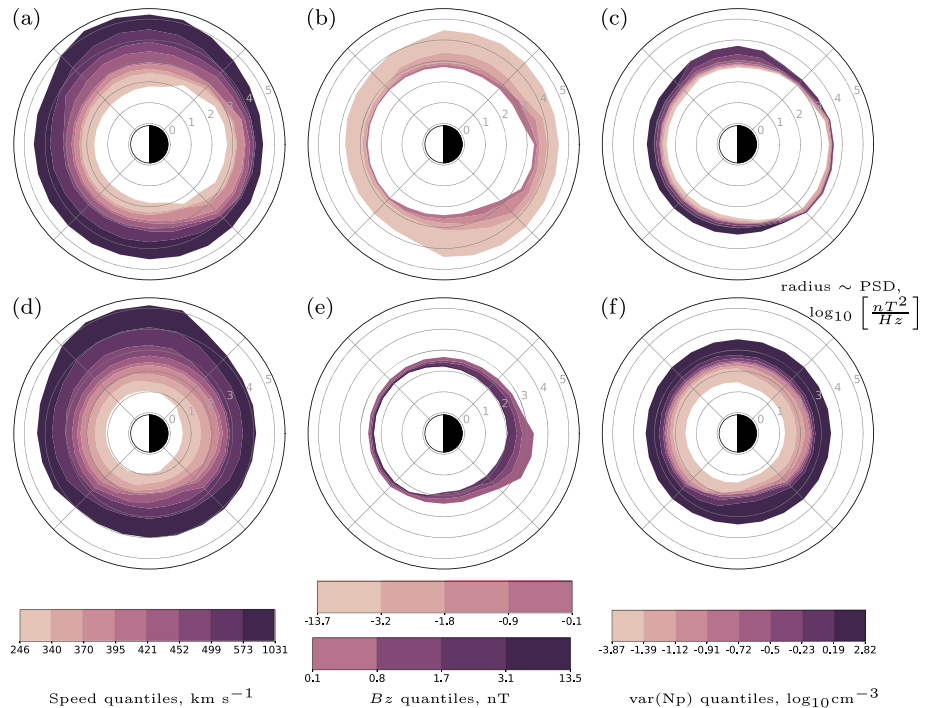


Figure 4. Variation of ULF wave power at GILL station, 5 mHz. Model-predicted power spectral density is shown by magnetic local time at quantiles of (a) speed (for median $B_z < 0$ and $\text{var}(Np)$), (b) $B_z < 0$ (for median speed and $\text{var}(Np)$), and (c) $\text{var}(Np)$ (for median speed and $B_z < 0$). Median values for speed, $B_z < 0$, and $\text{var}(Np)$ are 421 km s⁻¹, -1.8 nT, and $\text{var}(Np) = -0.716 \log_{10} (\text{cm}^{-3})$, respectively. (d-f) Also variation of wave power with speed, B_z , and $\text{var}(Np)$ but for $B_z > 0$ (with a median value of $B_z = 1.7$ nT held constant for (d) and (f)). Radius of each quantile corresponds to the power spectral density in $\log_{10} \left(\frac{\text{nT}^2}{\text{Hz}} \right)$ predicted for those solar wind values, at that station, frequency, and magnetic local time.

wave power on the dayside due to Kelvin-Helmholtz instabilities, which may be very difficult to distinguish from a more compressed magnetopause. Similarly, high dayside power with B_z is tied to other processes discussed above. As a result, the statement here (“more dayside power”) is weaker than the statement in the previous hypothesis (“more dayside power, especially at dawn”). Furthermore, because the magnetopause distance is only included indirectly via our parameter choice, the predictions from a random forest model are averaged over multiple magnetopause configurations, diluting any possible conclusions. We therefore choose not to make this a separate hypothesis as while magnetopause distance clearly has a significant impact on ULF wave power and ideally should be considered with other dominant dayside ULF processes, our model with the three parameters is not an appropriate tool to investigate that impact. Instead, we will simply consider the possible impact of magnetopause location when examining the other hypotheses above.

Hypothesis 3: The greatest remaining uncertainty will be in the midnight sector, with a possible additional increase in the dawn sector. Uncertainties will be larger for $B_z < 0$ than $B_z > 0$.

We expect that the greatest uncertainties in the model will occur in locations where our model less successfully captures the underlying physics due to our choice of parameters. Therefore, we expect this uncertainty to be largest in the midnight sector due to both the occurrence of substorm-generated power (which cannot be completely captured by solar wind properties) and the variable ground-to-L* correspondence, as stations will map to varying geomagnetic locations in the magnetotail, therefore averaging over different locations with varying properties. Similarly, there is likely to be larger uncertainty for $B_z < 0$ than $B_z > 0$ as we do not explicitly capture the effect of substorm behavior, and we expect more substorms for $B_z < 0$ (Freeman & Farrugia, 1999). Another possibility is the existence of uncertainty in regions where monochromatic standing waves develop: These will contribute significant amounts of power at a single frequency but are unlikely to be well represented by an average value over long time periods.

3.2. Testing the Hypotheses Using Model Power and Uncertainty

3.2.1. Hypothesis 1: Two B_z Regimes

In Figure 4 we show how ULF wave power occurrence varies in our model as we vary each parameter. For each plot, we hold two of the three parameters constant (at the median value) and vary the third, showing the predicted wave power between the minimum, maximum and quantiles. Figure 4a shows the power found in our model at GILL station, 5 mHz, as speed varies for $B_z < 0$ (i.e., showing power with quantiles of speed for median $B_z < 0$ and median $\text{var}(Np)$, $B_z = -1.8$ nT and $\text{var}(Np) = -0.716 \log_{10}(\text{cm}^{-3})$, respectively). Similarly, (b) shows the variation of power with quantiles of $B_z < 0$ for the median speed and $\text{var}(Np)$ (421 km s^{-1} and $\text{var}(Np) = -0.716 \log_{10}(\text{cm}^{-3})$). In (c) we vary $\text{var}(Np)$ for constant speed and $B_z < 0$, and (d)–(f) show the corresponding results for $B_z > 0$, with corresponding constants of 421 km s^{-1} , 1.7 nT, and $-0.716 \log_{10}(\text{cm}^{-3})$.

Figure 4 allows us to test Hypothesis 1. We do indeed see different behavior for $B_z \leq 0$ when varying both speed and $\text{var}(Np)$. For speed, the dawn-dusk asymmetry is clear for both $B_z \leq 0$ and there is a mid-night asymmetry for $B_z < 0$ (Figure 4a). Furthermore, we see that the $B_z = 0$ threshold indicates two distinct behaviour regimes by comparing Figure 4a to Figure 4d and Figure 4c to Figure 4f: For constant v_{sw} and $\text{var}(Np)$ values, respectively, the variation of power with B_z is very different for negative and positive B_z values. The $\text{var}(Np)$ plots are significantly different between $B_z \leq 0$, with $\text{var}(Np)$ accounting for very small variation of power across all MLT sectors for $B_z < 0$ and a much larger variation over all MLT where $B_z > 0$. These effects will be discussed further below.

We conclude that it is necessary to treat $B_z > 0$ and $B_z < 0$ as separate regimes and do so in the following analysis.

3.2.2. Hypotheses 2: Power Signature of External Drivers

We examine each solar wind parameter for evidence of a dawn, high-latitude increased power signature, postulated to indicate direct driving by magnetopause perturbations. To do so, we compare the change in each parameter at several stations for $B_z < 0$ (in Figure 5) and for $B_z > 0$ (in Figure 6). Both of these show the three parameters varying with quantile as before, for four stations FCHU, GILL, ISLL, and PINA.

We note that in the first column of both Figures 5 and 6 there is indeed a dawn-dusk asymmetry that is more pronounced in the higher-latitude stations, following increases with speed. This asymmetry is far less apparent with changes in B_z (second column) or $\text{var}(Np)$ (third column) even for similar amplitudes of wave power. This suggests that Hypothesis 2 is correct: that the dawn, high-latitude power signature may indeed represent a combination of driving by magnetopause oscillations and the different plasma density distribution in dawn, rather than simply the increased ability of the dawn sector to support waves. Otherwise, we would expect to see the same signature for other parameters, when total power is large enough to be comparable to power in the v_{sw} plots. We examine each subhypothesis below.

The dawn, high-latitude asymmetry is very clear for the speed column, suggesting that v_{sw} is indeed a good proxy for external perturbations driven by Kelvin-Helmholtz instabilities (KHI). This dawn asymmetry is only seen at lower-latitude stations during extremely fast solar wind speeds. However, it is also clear that even at lower-latitude stations and on the nightside, there is more power with increasing solar wind speed. These are not locations we would necessarily expect to contain ULF waves driven directly by KHI magnetopause perturbations, due to the density profile for low-latitude stations (and hence the Alfvén continuum Lee, 1996) and due to the magnetotail configuration and hence the open field lines and the varying distance to the magnetopause for the nightside. Our alternative hypothesis was that speed may work as a proxy for total power transferred from the solar wind to magnetosphere-scale perturbations and hence ULF wave power. We also discussed the complicating factor of high solar wind speed indicating a closer magnetopause. We now suggest that both these explanations may currently be indistinguishable from a less idealized view of KHI-driven wave power. In our idealized description, KHI-driven compressional waves travel radially and interact with the Alfvén continuum (i.e., the plasma density profile) in a relatively simple way. In a more realistic magnetosphere, magnetopause-driven waves propagate into the magnetosphere in all directions. The resulting “bouncing around” of the waves, particularly over the hourly windows considered here, may explain the global increases of power as an increases in solar wind speed. However, we would intuitively still expect less variation of power with speed at lower latitudes and at midnight.

Instead, the remaining power observed in Figure 6 at these locations suggests that there may be a more complicated relationship with speed, in addition to more realistic driving by KHI. Part of this complicated

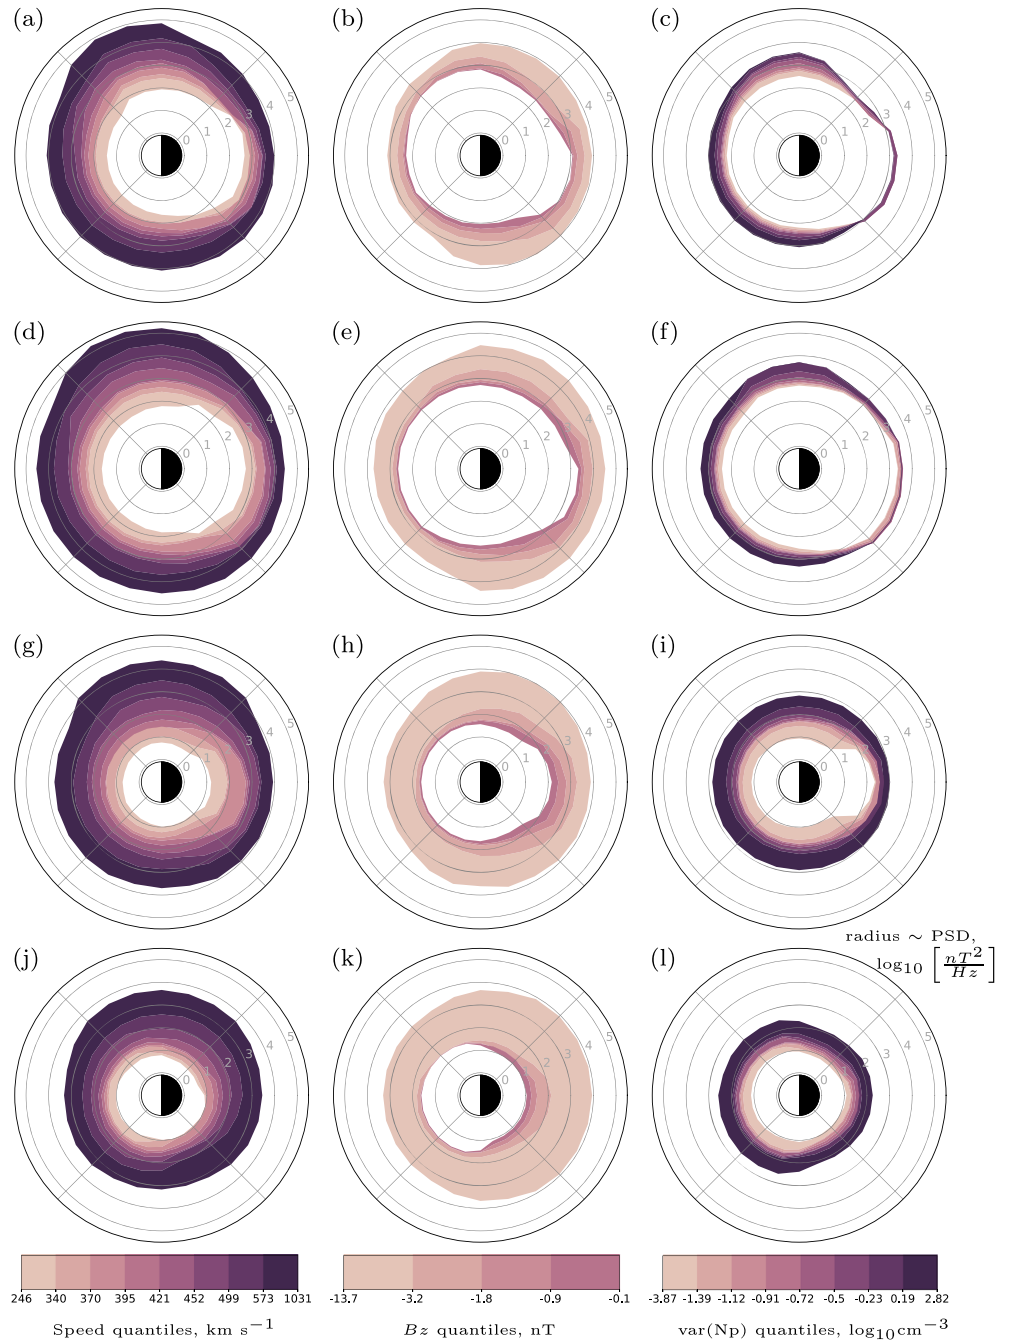


Figure 5. Variation of power with three solar wind parameters in the decision tree ensemble models at 5 mHz, by magnetic local time, for $B_z < 0$. By row: (a–c) FCHU ($L \sim 7.94$), (d–f) GILL ($L \sim 6.51$), (g–i) ISLL ($L \sim 5.40$), and (j–l) PINA ($L \sim 4.21$). By column: variation of power with speed, $B_z < 0$, and $\text{var}(Np)$. Radius corresponds to predicted power spectral density.

relationship likely includes the effect of v_{sw} representing a closer magnetopause, which, as we have discussed, is not easily investigated with our model. We therefore suggest that Hypothesis 2a is satisfied, but that there are further nuances to the distribution of energy throughout the magnetosphere during times of increasing solar wind speed.

We can also address the cause of the B_z driving. In Hypothesis 2b, we suggested that with increasingly negative $B_z < 0$ we might observe either a dawn-dusk asymmetry in wave power arising from magnetopause deformations by traveling flux tubes or a midnight sector enhancement associated with substorms. These




Figure 6. Variation of power with three solar wind parameters in the decision tree ensemble models at 5 mHz, by magnetic local time, for $B_z > 0$. By row: (a–c) FCHU ($L \sim 7.94$), (d–f) GILL ($L \sim 6.51$), (g–i) ISLL ($L \sim 5.40$), and (j–l) PINA ($L \sim 4.21$). By column: variation of power with speed, $B_z < 0$, and $\text{var}(Np)$. Radius corresponds to predicted power spectral density.

options can be tested using the $B_z \leq 0$ split between Figures 5 and 6, in addition to the variation of power with all B_z in the second column.

Comparing the first (i.e., speed) columns of both Figure 5 ($B_z < 0$) and Figure 6 ($B_z > 0$), we find that the dawn-dusk asymmetry seen with increases in speed appears modestly larger for $B_z < 0$. There is little evidence of a dawn, high-latitude increase in power with variations of B_z (second column of Figures 5 and 6), suggesting that B_z represents minimal driving by magnetopause perturbations.

On the other hand, there is a premidnight increase in power for all except the most strongly positive values of B_z (second column of Figures 5 and 6). This premidnight increase is strong enough to also be reflected with variation of speed for $B_z < 0$ (first column, Figure 5) and is centered at 23 MLT, the most common location for substorms (Gjerloev et al., 2004). The increase in power is particularly clear for auroral stations FCHU and GILL, with a smaller effect for the lower-latitude stations which is centered closer to midnight.

We note that in the $B_z < 0$ column of Figure 5 there is residual increased power focused on the nightside but observed at all MLTs, whereas in contrast there appears to be very little variation of power with $B_z > 0$, except for values of B_z close to 0 which likely still include substorm activity (Freeman & Farrugia, 1999). This residual power with $B_z < 0$ is probably due to substorm activity driving global ULF wave propagation on our hourly timescale as it is seen at all MLTs but could also be related to other dayside processes associated with $B_z < 0$ such as a closer magnetopause. These processes cannot be distinguished at this time, and so we cannot comment on their contribution.

Together, the lack of external power signature and the increase in power at 23 MLT indicate that power increases with negative B_z are unlikely to represent significant direct driving by magnetopause perturbations. Instead, power variations with B_z appear to indicate that substorm activity is the dominant process of the two

To investigate the role of $\text{var}(Np)$ as a proxy for external magnetopause perturbations, we intended to examine the variation of power with $\text{var}(Np)$ in the third column of both Figures 5 and 6. However, we note that the third column of Figure 5 looks quite unexpected, with barely any variation of power with solar wind $\text{var}(Np)$. This unexpected behavior could be physical, an artifact of our modeling method, or a combination of the two.

A physical interpretation of the lack of power variation with $\text{var}(Np)$ at high latitude in midnight may be due to those stations predominantly connecting to open field lines while $B_z < 0$. It is not clear why this would impact the $\text{var}(Np)$ column so drastically and the speed column not at all. On the other hand, these artifacts could be due to the interdependence of the parameters and the operation of the decision trees. There may have been fewer splits made along $\text{var}(Np)$ because significantly more variation existed along other parameters, or because $\text{var}(Np)$ covaries particularly efficiently with v_{sw} and/or B_z under these conditions. In both these cases the model would be working as expected to accurately approximate the variation of power in the solar wind parameter space; unfortunately, interparameter relationships make it difficult to analyze the physics behind variation of power in the model. Given that these low-variation situations are in the midnight sector, for high-latitude stations when $B_z < 0$, we think it most likely that the cause is predominantly physical (i.e., open field lines) but complicated by solar wind interparameter relationships. We choose to only draw conclusions from Figure 6f, 6i, and 6l for Hypothesis 2c.

For $\text{var}(Np)$, we do not see any evidence of the power signature associated with magnetopause perturbations. Excluding those figures whose results are contaminated by open field lines, power appears symmetric with MLT, or with an increase on the entire dayside for the lowest latitude stations. We conclude that $\text{var}(Np)$ does not represent power directly driven by waves from the magnetopause. This is consistent with previous work suggesting that short-lived solar wind processes such as impulsive density perturbations are screened out by the magnetopause, which cannot respond fast enough and therefore acts as a low-pass filter (e.g., Archer et al., 2013). Instead, any increase in power must be due to another aspect of sudden global compressions, and so this is a question which should be addressed in future work.

In this second hypothesis we have essentially been investigating the dawn-dusk asymmetry. However, we have not yet discussed the impact of asymmetries in ionospheric conductivity. This affects the transmission of ULF waves through the ionosphere and therefore modulates wave observation by ground magnetometers. (Glassmeier & Stellmacher, 2000) found a greater dawn-dusk asymmetry for ground observations than for in situ observations, suggesting that differing ionospheric properties contribute to the observed wave power asymmetry. However, we do not believe that this affects the results of our analysis. The ground station used in Glassmeier & Stellmacher (2000) is on a similar geomagnetic latitude to GILL (second row of Figures 5 and 6) whereas their in situ data corresponds to ground station PINA (bottom row of Figures 5 and 6) which shows significantly less dawn-dusk asymmetry, similar to the in situ results. These more closely located measurements indicate that the difference between ground and in situ data may be predominantly

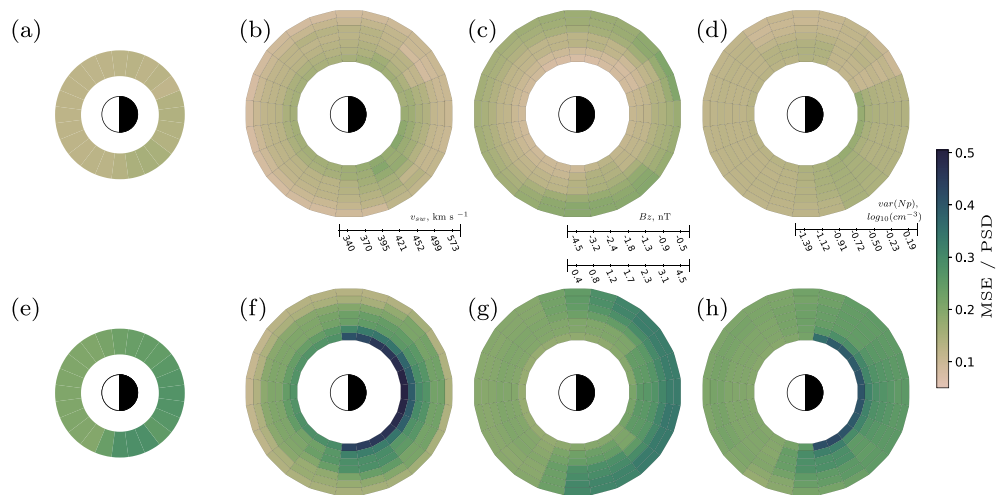


Figure 7. GILL 5.0 mHz. Top row is for $B_z < 0$, bottom for $B_z > 0$. From left to right: (a) uncertainty in each MLT bin for median speed, B_z , and $\text{var}(N_p)$ values (421 km s^{-1} , -1.8 and 1.7 nT , and $-0.716 \log_{10} (\text{cm}^{-3})$, respectively). (b) The remaining uncertainty at speed quantiles for each MLT (quantile values shown in the bar in the bottom right corner), (c) similar uncertainty for $B_z < 0$ at median values of speed and $\text{var}(N_p)$, (d) uncertainty for $\text{var}(N_p)$ by MLT. (e–h) The same for $B_z > 0$.

due to latitudinal differences, rather than ionospheric processing. Furthermore, if the dawn-dusk asymmetry were predominantly due to asymmetries in ionospheric conductivity we would see the asymmetry for more parameters and latitudes than in our analysis above. However, the relationship between ionospheric conductivity and solar wind properties (e.g., Newell et al., 2009) and the related effect of the auroral oval on observed wave power (Kozyreva et al., 2016; Pilipenko et al., 2001; Simms et al., 2006) are likely to have additional effects that should be investigated in future, perhaps with a secondary set of hypotheses, or with an expanded set of parameters if they could be causally established.

3.2.3. Hypothesis 3: Uncertainty in Our Model

We also investigate azimuthal asymmetries by examining which MLTs contain the most remaining uncertainty. Large amounts of remaining uncertainty indicate areas where we approximate the physics less well, due to either our choice of parameters or model type. As each decision tree is trained on MSE (and this quantity is stored in the final tree), we use a MSE-based quantity to determine uncertainty. Instead of pure MSE, which would simply scale with total power, we characterize the remaining uncertainty as the MSE divided by the predicted power. In this way a remaining uncertainty of ~ 1 would indicate that the uncertainty is comparable to the power predicted in that bin. To calculate this remaining uncertainty for a given point in the parameter space, we identify the appropriate bin in each decision tree in the ensemble and fetch the MSE associated with that bin. Thus the remaining uncertainty is average MSE remaining in all trees of the ensemble at that point, divided by the power predicted by the ensemble (i.e., the average power predicted by trees in the ensemble).

Figure 7 shows the uncertainty in decision tree predictions at selected points in the parameter space. From left to right: First we show uncertainty across all decision trees in an ensemble for each MLT at the median values of speed, $B_z < 0$ and $\text{var}(N_p)$. Then we show the uncertainty at each MLT for octiles of speed, for the specified value median values of $B_z < 0$ and $\text{var}(N_p)$. Outer quantiles—the minimum and maximum values of speed—are excluded here, as they are the most poorly represented regions of the parameter space and therefore the remaining uncertainty is less physically informative. Similarly, we show variation of uncertainty for each MLT with $B_z < 0$ (in (c)) and $\text{var}(N_p)$ (in (d)). The bottom row is the same but for $B_z > 0$. Each bin shown in Figure 7 represents one point in the parameter space; neighboring points may well still end up in one leaf node of the tree and therefore have the same MSE as they are in the same decision tree bin.

We suggested that the greatest remaining uncertainty would be in midnight based on our assumption that substorms were the most poorly approximated physics in our model. While the midnight sector is indeed the MLT sector with the most uncertainty, surprisingly there is more uncertainty for all parameters for $B_z > 0$. This result suggests that we are approximating the physics better for $B_z < 0$ than $B_z > 0$. This can still be

understood in terms of substorms, if one takes them to be the dominant ULF-effective B_z -related process distinguishable in our model (which they appear to be, based on the Hypothesis 2b analysis).

If $B_z < 0$ is a good proxy for substorms, it would represent large impacts from substorms, with only a small amount of variability. Conversely, for $B_z > 0$, we may observe significant substorm activity, or none at all (see, e.g., Freeman & Farrugia, 1999; Lee et al., 2010; Wygant et al., 1983). In this case $B_z > 0$ may on average represent low amounts of substorm activity, but with high variability. If ULF wave power is related to substorm activity, the poor quality of $B_z > 0$ as a substorm predictor would be reflected in the uncertainty here. For example, it has been shown that the amount of tail loading (and hence substorm occurrence) depends on the time since the interplanetary magnetic field was directed southward (Wygant et al., 1983). Reducing remaining substorm-related uncertainty would require a better proxy for the substorm cycle. We note that this B_z -dependent uncertainty is seen at all MLT, and not solely in the midnight sector. This uncertainty in dayside wave power may be due to the timescale of our observations, as the effect of substorms may well propagate to the dayside within an hour. However, the uncertainty may additionally be caused by other B_z -dependent effects such as those described when constructing Hypothesis 2b (magnetopause location, Alfvén speed, Kelvin-Helmholtz stability of magnetopause and also ionospheric conductivity). Including a substorm activity proxy would allow this to be tested, and may improve dayside power estimates.

We also note that there appears to be little dawn-dusk asymmetry in the remaining uncertainty, suggesting that capturing the source of such an asymmetry is not the most significant remaining problem in the characterization of ULF wave power.

Overall the largest amount of uncertainty is for $B_z > 0$, at low speed and low $\text{var}(Np)$ values. These are the conditions under which there is the least ULF wave driving by solar wind properties. As the uncertainty here is significantly larger, we suggest that power is poorly parameterized by our solar wind properties because external driving no longer dominates so completely; instead, the internal state of the magnetosphere (including the equilibrium magnetopause location) and internal drivers may be more significant. This should be confirmed in future work as it indicates that internal magnetospheric properties may be a significant second component of more accurate ULF wave models.

3.2.4. Summary

In this section we tested hypotheses about ULF wave power occurrence in MLT. We found it necessary to split the analysis into two separate regimes ($B_z \leq 0$). We found that a dawn asymmetry of higher power with a peak at high latitudes can only be found with increasing speed, and is not observed for other parameters even when the total power is comparable. As speed is the parameter most strongly associated with direct driving by magnetopause oscillations, this finding suggests that the dawn asymmetry is a combined result of both direct driving by magnetopause perturbations and the more hospitable density profile at dawn, rather than solely the ability of the dawn sector to support ULF waves. We conclude that the power increase with $B_z < 0$ is more likely to be due to substorms than to external drivers. We are unable to fully analyze variation of power with $\text{var}(Np)$, but conclude that it does not represent driving by magnetopause perturbations.

Remaining uncertainty in the model suggested that we capture substorms poorly, although contrary to expectations, this effect was greater for $B_z > 0$ and was not restricted to the midnight sector. Finally, the most remaining uncertainty was found where the solar wind driving was least strong. This result suggests that the internal state of the magnetosphere is comparable to the smallest amounts of solar wind driving and may therefore have a significant contribution to ULF wave power that should be included in future models.

4. Discussion

In testing, the decision tree ensemble models perform modestly better than the previous empirical models at predicting real data while also presenting significant improvements in methodology and utility. They are more robust as the bin choices are not arbitrary, and remaining error can be reduced using methods such as parameter optimization and ensemble averaging. They are also easier to train, save, use and share as part of the Python module scikit-learn. They have better coverage as, with no empty bins, predictions can be made across the entire parameter space. Finally, the variable bin size means that, by definition, they capture the areas of parameter space with more rapidly changing ULF wave power, such as around $B_z \sim 0$ and by MLT. While characterization of the full spectrum of wave behavior would require properties that cannot be captured in a ground-based model (e.g., the wave mode and azimuthal wave structure, or m-number), the high resolution empirical models (and their construction method) presented in this paper represent

a significantly increased capability to investigate global magnetospheric wave processes. For example, in section 3 existing theories of ULF wave generation and propagation in the parameter space were tested with MLT.

Our use of decision trees is based on several underlying difficulties in analyzing space physics data such as an unevenly sampled parameter space and the interdependence of driving parameters—in particular, the effect of nonlinear relationships between those parameters. Each of these is addressed somewhat by choosing to use a decision tree.

In the decision tree algorithm, bins are determined by the variance in the data across the parameter space. As a result, the algorithm is likely to split large bins that are heavily populated at one end but represent different behavior across the bin. Therefore, this algorithm immediately reduces (although does not remove) the bias from such bins in comparison to large linearly spaced bins. The decision tree algorithm allows a higher bin resolution where we have more data but also prevents bins calculating a mean power from only three or four points. It is possible the outer edge of the parameter space is still poorly characterized with sparsely populated bins; this is unavoidable and would improve with more data.

The driving parameter interdependence has two aspects. First, interpretability. The parameters used in this study were chosen as they were found to be causally correlated to changes in ULF wave power (Bentley et al., 2018). This removes parameters correlated to but not driving ULF wave power. Therefore, for our decision tree ensembles, the interpretability and interdependence of our parameters depends entirely on the choice of $[v_{sw}, Bz, \text{var}(Np)]$ from previous studies. Second, the effect of bias. This results from the inherent capture of interparameter relationships in our machine learning model; we are essentially capturing the relationship between solar wind parameters which is due to the nature of solar wind generation and propagation (Borovsky, 2018), whereas ideally we only want the effect of the parameters on the driving of ULF waves. This complication can be explained succinctly by using two functions to represent the solar wind driving of ULF waves power P approximated with our model:

$$\begin{aligned} g : \text{Sun} &\rightarrow [v_{sw}, Bz, \text{var}(Np)]^a \\ f : [v_{sw}, Bz, \text{var}(Np), \text{freq}, \text{lat}, \text{component}]^b &\rightarrow P. \end{aligned} \quad (2)$$

Generation and propagation of the solar wind is determined by the Sun (function g). However, this means that the values of $[v_{sw}, Bz, \text{var}(Np)]^a$ which are observed in the solar wind near Earth do not include all possible combinations of the parameters $v_{sw}, Bz, \text{var}(Np)$. The model (f) is designed to take all possible permutations in this parameter space, and intuitively we expect the output to represent only the relationship between solar wind variables observed near the Earth and the ULF wave power. Instead, the input to our random forest model (f) predicting ULF wave PSD P is actually the subset $[v_{sw}, Bz, \text{var}(Np)]^a$ of all the possible $[v_{sw}, Bz, \text{var}(Np)]$ values in our parameter space. As we do not explicitly remove or account for the effect of g on our available input choices, $[v_{sw}, Bz, \text{var}(Np)]^b = [v_{sw}, Bz, \text{var}(Np)]^a$, and so these unknown interparameter relationships are inherently included in our model function f instead. As a result, there is a bias toward values of power seen under more common solar wind conditions, due to the indirect inclusion of solar wind processing in our model.

To a large extent this relationship is unavoidable. However, we can somewhat reduce the effect of these inherently included interparameter relationships by using a decision tree. We do not assume that relationships between input parameters, or between input parameters and power P are linear, or monotonic, or a single function. Instead, the total combined function f is approximated using a series of piecewise step functions across the parameter space. In this way the nonlinear relationship between power and Bz was found to be more significant than anticipated; the two regimes either side of $Bz = 0$ behave very differently. However, that is not to say that we have completely removed the effect of the interparameter relationships. Instead, we have only approximated the final (nonlinear) combined function representing the effect of our parameters on power. When the decision tree returns the correct power value for a given set of solar wind conditions, it is not necessarily because the contribution of each parameter has been independently captured but may instead be tied up with interparameter relationships inherently approximated by the model. In this case the model will perform well for forecasting and predictions but is more difficult to interpret, such as our attempt to analyze the effect of $\text{var}(Np)$ on wave power and our inability to include magnetopause distance in our first set of hypotheses.

The effectivity of the decision tree model also suggests that more parameters could be added. The model here has six dimensions and is already difficult to analyze, necessitating a hypothesis testing rather than a statistical study approach. However, including more parameters and testing more hypotheses could yield greater insight into the physics. For example, we could include F10.7 to investigate the solar cycle (Murphy et al., 2011) or parameters such as estimated magnetopause location or internal properties of the magnetosphere (Rae et al., 2019). Parameter interdependence would be an important consideration for such additions. Other parameter choices will be investigated in future.

5. Conclusion

In this article we have constructed a series of decision tree ensembles characterizing ULF wave PSD from ground-based magnetometers, driven by solar wind speed, B_z and variance of proton number density. When trained correctly, decision trees ensembles are less susceptible to bias and variance than simple binned statistical models as they choose bins intelligently in the parameter space. As a result we can predict how much power will occur more accurately and with far greater resolution in the input parameters, such as in MLT. Models span 1–15 mHz and cover roughly $L \sim 4$ to ~ 8 . Using forecasting skill, we find that this model outperforms the previous model Bentley et al. (2019) and is more versatile.

We have proposed that testing a series of hypotheses is a suitable method for investigating successively finer contributions in parameterized models such as the one here. As an example we have analyzed MLT variations in our six-dimensional model using this hypothesis testing, rather than a full statistical survey, for a single representative frequency 5 mHz. We confirmed that for ULF wave occurrence, $B_z \leq 0$ represents two different physical regimes due to substorm activity. We have also confirmed that external driving by solar wind speed via Kelvin-Helmholtz instabilities is seen in our model and that a high power signature in dawn at high latitudes corresponds to the combined effect of directly driven magnetopause perturbations and a shallower radial plasma density profile. We used model uncertainty to identify where the physics was poorly approximated; while large amounts of uncertainty do appear to be due to substorm activity, this appears more for $B_z > 0$ than $B_z < 0$ and is seen globally rather than just in the midnight sector. Furthermore, the largest uncertainty is seen for $B_z > 0$ and low speed or $\text{var}(N_p)$. This indicates that while solar wind driving dominates the ULF wave power observed, the internal state of the magnetosphere is likely to have a significant secondary role.

Future work on this topic includes creating probabilistic radial diffusion coefficients based on the model used here and the inclusion of further parameters both to understand the underlying physics and for forecasting purposes.

Data Availability Statement

Data from CARISMA can be found online (at <https://carisma.ca>). OMNI data are available from the GSFC/SPDF OMNIWeb interface (at <https://omniweb.gsfc.nasa.gov/>). The decision tree ensembles are available to use from Zenodo (Bentley et al., 2020).

References

- Archer, M. O., Horbury, T. S., Eastwood, J. P., Weygand, J. M., & Yeoman, T. K. (2013). Magnetospheric response to magnetosheath pressure pulses: A low-pass filter effect. *Journal of Geophysical Research: Space Physics*, 118, 5454–5466. <https://doi.org/10.1002/jgra.50519>
- Baker, G. J., Donovan, E. F., & Jackel, B. J. (2003). A comprehensive survey of auroral latitude Pc5 pulsation characteristics. *Journal of Geophysical Research*, 108(A10), 1384. <https://doi.org/10.1029/2002JA009801>
- Bentley, S. N. (2019). Predicting the effect of the solar wind on the ultra-low frequency plasma waves driving Earth's radiation belts. (PhD Thesis), University of Reading. <https://doi.org/10.17864/1926.85255>
- Bentley, S. N., Stout, J., Bloch, T., & Watt, C. E. J. (2020). Dataset: Random forest models of magnetospheric ultra-low frequency wave power. <https://doi.org/10.5281/zenodo.3828506>
- Bentley, S. N., Watt, C. E. J., Owens, M. J., & Rae, I. J. (2018). ULF wave activity in the magnetosphere: Resolving solar wind interdependencies to identify driving mechanisms. *Journal of Geophysical Research: Space Physics*, 123, 2745–2771. <https://doi.org/10.1002/2017JA024740>
- Bentley, S. N., Watt, C. E. J., Rae, I. J., Owens, M. J., Murphy, K., Lockwood, M., & Sandhu, J. K. (2019). Capturing uncertainty in magnetospheric ultralow frequency wave models. *Space Weather*, 17, 599–618. <https://doi.org/10.1029/2018SW002102>
- Berner, J., Achatz, U., Batté, L., Bengtsson, L., De La Cámara, A., Christensen, H. M., et al. (2017). Stochastic parameterization toward a new view of weather and climate models. *Bulletin of the American Meteorological Society*, 98(3), 565–587. <https://doi.org/10.1175/BAMS-D-15-00268.1>
- Berube, D., Sanny, J., Taus, R., & Garoutte, A. (2014). Dayside distribution of Pc5 wave power in the quiet magnetosphere and its response to the solar wind. *Planetary and Space Science*, 97, 1–8. <https://doi.org/10.1016/j.pss.2014.04.012>

Acknowledgments

The authors thank I. R. Mann, D. K. Milling, and the rest of the CARISMA team for data. CARISMA is operated by the University of Alberta, funded by the Canadian Space Agency. The authors gratefully acknowledge low-resolution OMNI and would also like to extend their thanks for Python scientific computing packages used in this project including `scikit-learn` (Grisel et al., 2020; Pedregosa et al., 2011), `pandas` (McKinney, 2010; Reback et al., 2019), `seaborn` (Waskom et al., 2018), and `xarray` (Hoyer et al., 2019). Decision tree ensembles were saved using `joblib` (<https://joblib.readthedocs.io/>) and `sklearn-json` (<https://github.com/saromanov/scikit-json>). This work was supported by STFC Grant ST/R000921/1 and NERC Grant NE/P017274/1. T. B. was supported by STFC doctoral training Grant ST/R505031/1, and J. S. was supported by the University of Reading UROP scheme.

- Borovsky, J. E. (2018). On the origins of the intercorrelations between solar wind variables. *Journal of Geophysical Research: Space Physics*, 123, 20–29. <https://doi.org/10.1002/2017JA024650>
- Claudepierre, S. G., Hudson, M. K., Lotko, W., Lyon, J. G., & Denton, R. E. (2010). Solar wind driving of magnetospheric ULF waves: Field line resonances driven by dynamic pressure fluctuations. *Journal of Geophysical Research*, 115, A11202. <https://doi.org/10.1029/2010JA015399>
- Elkington, S. R., Hudson, M. K., & Chan, A. A. (1999). Acceleration of relativistic electrons via drift-resonant interaction with toroidal-mode Pc-5 ULF oscillations. *Geophysical Research Letters*, 26(21), 3273–3276. <https://doi.org/10.1029/1999GL003659>
- Freeman, M. P., & Farrugia, C. J. (1999). Solar wind input between substorm onsets during and after the October 18–20, 1995, magnetic cloud. *Journal of Geophysical Research*, 104(A10), 22,729–22,744. <https://doi.org/10.1029/1999JA900204>
- Gjerloev, J. W., Hoffman, R. A., Friel, M. M., Frank, L. A., & Sigwarth, J. B. (2004). Substorm behavior of the auroral electrojet indices. *Annales Geophysicae*, 22(6), 2135–2149. <https://doi.org/10.5194/angeo-22-2135-2004>
- Glassmeier, K.-H., & Stellmacher, M. (2000). Concerning the local time asymmetry of Pc5 wave power at the ground and field line resonance widths. *Journal of Geophysical Research*, 105(A8), 18,847–18,855. <https://doi.org/10.1029/2000JA900037>
- Glauert, S. A., Horne, R. B., & Meredith, N. P. (2014). Three-dimensional electron radiation belt simulations using the BAS Radiation Belt Model with new diffusion models for chorus, plasmaspheric hiss, and lightning-generated whistlers. *Journal of Geophysical Research: Space Physics*, 119, 268–289. <https://doi.org/10.1002/2013JA019281>
- Grisel, O., Mueller, A., Lars, Gramfort, A., Louppe, G., Prettenhofer, P., et al. (2020). Scikit-learn/scikit-learn: Scikit-learn 0.22.1. <https://doi.org/10.5281/ZENODO.3596890>
- Hartinger, M. D., Turner, D. L., Plaschke, F., Angelopoulos, V., & Singer, H. (2013). The role of transient ion foreshock phenomena in driving Pc5 ULF wave activity. *Journal of Geophysical Research: Space Physics*, 118, 299–312. <https://doi.org/10.1029/2012JA018349>
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning*. New York, NY: Springer New York. <https://doi.org/10.1007/b94608>
- Horne, R. B., Glauert, S. A., Meredith, N. P., Boscher, D., Maget, V., Heynderickx, D., & Pitchford, D. (2013). Space weather impacts on satellites and forecasting the Earth's electron radiation belts with SPACECAST. *Space Weather*, 11, 169–186. <https://doi.org/10.1002/swe.20023>
- Hoyer, S., Hamman, J., Roos, M., Fitzgerald, C., Cherian, D., Maussion, F., et al. (2019). pydata/xarray v0.14.1. <https://doi.org/10.5281/ZENODO.3547868>
- Hwang, K. J., & Sibeck, D. G. (2016). Role of low-frequency boundary waves in the dynamics of the dayside magnetopause and the inner magnetosphere. In *Low-frequency waves in space plasmas* (pp. 213–239). <https://doi.org/10.1002/9781119055006.ch13>
- King, J. H., & Papitashvili, N. E. (2005). Solar wind spatial scales in and comparisons of hourly Wind and ACE plasma and magnetic field data. *Journal of Geophysical Research*, 110, A02104. <https://doi.org/10.1029/2004JA010649>
- Kozyreva, O., Pilipenko, V., Engebretson, M., Klimushkin, D., & Mager, P. (2016). Correspondence between the ULF wave power spatial distribution and auroral oval boundaries. *Solnechno-Zemnaya Fizika*, 2(2), 35–45. <https://doi.org/10.12737/16848>
- Langan, R., Archibald, R., Plumlee, M., Mahajan, S., Ricciuto, D., Yang, C., et al. (2014). Stochastic parameterization to represent variability and extremes in climate modeling. *Procedia Computer Science*, 29, 1146–1155. <https://doi.org/10.1016/j.procs.2014.05.103>
- Lee, D.-H. (1996). Dynamics of MHD wave propagation in the low-latitude magnetosphere. *Journal of Geophysical Research*, 101(A7), 15,371–15,386. <https://doi.org/10.1029/96JA00608>
- Lee, D.-Y., Choi, K.-C., Ohtani, S., Lee, J. H., Kim, K. C., Park, K. S., & Kim, K.-H. (2010). Can intense substorms occur under northward IMF conditions? *Journal of Geophysical Research*, 115, A01211. <https://doi.org/10.1029/2009JA014480>
- Leutbecher, M., Lock, S.-J., Ollinaho, P., Lang, S. T. K., Balsamo, G., Bechtold, P., et al. (2017). Stochastic representations of model uncertainties at ECMWF: State of the art and future vision. *Quarterly Journal of the Royal Meteorological Society*, 143(707), 2315–2339. <https://doi.org/10.1002/qj.3094>
- Ma, Q., Li, W., Bortnik, J., Thorne, R. M., Chu, X., Ozeke, L. G., et al. (2018). Quantitative evaluation of radial diffusion and local acceleration processes during GEM challenge events. *Journal of Geophysical Research: Space Physics*, 123, 1938–1952. <https://doi.org/10.1002/2017JA025114>
- Mann, I. R., Milling, D. K., Rae, I. J., Ozeke, L. G., Kale, A., Kale, Z. C., et al. (2008). The Upgraded CARISMA Magnetometer Array in the THEMIS Era. *Space Science Reviews*, 141(1–4), 413–451. <https://dx.doi.org/10.1007/s11214-008-9457-6>
- Mann, I. R., Murphy, K. R., Ozeke, L. G., Rae, I. J., Milling, D. K., & Kale, A. (2012). The role of ultralow frequency waves in radiation belt dynamics. *Geophysical Monograph Series*, 199, 69–91. <https://doi.org/10.1029/2012GM001349>
- Mann, I. R., & Wright, A. N. (1999). Diagnosing the excitation mechanisms of Pc5 magnetospheric flank waveguide modes and FLRs. *Geophysical Research Letters*, 26(16), 2609–2612. <https://doi.org/10.1029/1999GL900573>
- Mathie, R., & Mann, I. R. (2001). On the solar wind control of Pc5 ULF pulsation power at mid-latitudes: Implications for MeV electron acceleration in the outer radiation belt. *Journal of Geophysical Research*, 106(A12), 29,783–29,796. <https://doi.org/10.1029/2001JA000002>
- McKinney, W. (2010). Data structures for statistical computing in Python. In S. van der Walt & J. Millman (Eds.), *Proceedings of the 9th Python in Science Conference* (pp. 56–61). <https://doi.org/10.25080/Majora-92b1922-00a>
- Mehta, P., Bukov, M., Wang, C.-h., Day, A. G., Richardson, C., Fisher, C. K., & Schwab, D. J. (2019). A high-bias, low-variance introduction to machine learning for physicists. *Physics Reports*, 810, 1–124. <https://doi.org/10.1016/j.physrep.2019.03.001>
- Murphy, K. R., Mann, I. R., Jonathan Rae, I., & Milling, D. K. (2011). Dependence of ground-based Pc5 ULF wave power on F10.7 solar radio flux and solar cycle phase. *Journal of Atmospheric and Solar-Terrestrial Physics*, 73(11–12), 1500–1510. <https://doi.org/10.1016/j.jastp.2011.02.018>
- Murphy, K. R., Mann, I. R., & Sibeck, D. G. (2015). On the dependence of storm time ULF wave power on magnetopause location: Impacts for ULF wave radial diffusion. *Geophysical Research Letters*, 42, 9676–9684. <https://doi.org/10.1002/2015GL066592>
- Newell, P. T., Sotirelis, T., & Wing, S. (2009). Diffuse, monoenergetic, and broadband aurora: The global precipitation budget. *Journal of Geophysical Research*, 114, A09207. <https://doi.org/10.1029/2009JA014326>
- Oshiro, T. M., Perez, P. S., & Baranauskas, J. A. (2012). How many trees in a random forest? In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, LNAI (Vol. 7376, pp. 154–168). https://doi.org/10.1007/978-3-642-31537-4_13
- Pahud, D. M., Rae, I. J., Mann, I. R., Murphy, K. R., & Amalraj, V. (2009). Ground-based Pc5 ULF wave power: Solar wind speed and MLT dependence. *Journal of Atmospheric and Solar-Terrestrial Physics*, 71(10–11), 1082–1092. <https://doi.org/10.1016/j.jastp.2008.12.004>
- Pedregosa, F., Varoquax, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825–2830.

- Pilipenko, V. A., Watermann, J., Popov, V. A., & Papitashvili, V. O. (2001). Relationship between auroral electrojet and Pc5 ULF waves. *Journal of Atmospheric and Solar-Terrestrial Physics*, 63(14), 1545–1557. [https://doi.org/10.1016/S1364-6826\(01\)00031-1](https://doi.org/10.1016/S1364-6826(01)00031-1)
- Rae, I. J. (2017). Dawn-dusk asymmetries in ultra-low-frequency waves. In S. Haaland, A. Runov, & C. Forsyth (Eds.), *Dawn-dusk asymmetries in planetary plasma environments*. <https://doi.org/10.1002/9781119216346.ch14>
- Rae, I. J., Mann, I. R., Murphy, K. R., Ozeke, L. G., Milling, D. K., Chan, A. A., et al. (2012). Ground-based magnetometer determination of in situ Pc4–5 ULF electric field wave spectra as a function of solar wind speed. *Journal of Geophysical Research*, 117, A04221. <https://doi.org/10.1029/2011JA017335>
- Rae, I. J., Murphy, K. R., Watt, C. E. J., Sandhu, J. K., Georgiou, M., Degeling, A. W., et al. (2019). How do ultra low frequency waves access the inner magnetosphere during geomagnetic storms? *Geophysical Research Letters*, 46, 10,699–10,709. <https://doi.org/10.1029/2019GL082395>
- Reback, J., McKinney, W., den Bossche, J. V., Jbrockmendel, Augspurger, T., Cloud, P., et al. (2019). pandas-dev/pandas: v0.25.3. <https://doi.org/10.5281/ZENODO.3524604>
- Reeves, G. D., Chen, Y., Cunningham, G. S., Friedel, R. W. H., Henderson, M. G., Jordanova, V. K., et al. (2012). Dynamic Radiation Environment Assimilation Model: DREAM. *Space Weather*, 10, S03006. <https://doi.org/10.1029/2011SW000729>
- Rostoker, G., Samson, J. C., Creutzberg, F., Hughes, T. J., McDiarmid, D. R., McNamara, A. G., et al. (1995). Canopus—A ground-based instrument array for remote sensing the high latitude ionosphere during the ISTEP/GGS program. *Space Science Reviews*, 71(1–4), 743–760. <https://doi.org/10.1007/BF00751349>
- Shue, J.-H., Song, P., Russell, C. T., Steinberg, J. T., Chao, J. K., Zastenker, G., et al. (1998). Magnetopause location under extreme solar wind conditions. *Journal of Geophysical Research*, 103(A8), 17,691–17,700. <https://doi.org/10.1029/98JA01103>
- Simms, L. E., Engebretson, M. J., Posch, J. L., & Hughes, W. J. (2006). Effects of the equatorward auroral boundary location and solar wind parameters on Pc5 activity at auroral zone stations: A multiple regression analysis. *Journal of Geophysical Research*, 111, A10217. <https://doi.org/10.1029/2005JA011587>
- Simms, L. E., Pilipenko, V. A., & Engebretson, M. J. (2010). Determining the key drivers of magnetospheric Pc5 wave power. *Journal of Geophysical Research*, 115, A10241. <https://doi.org/10.1029/2009JA015025>
- Subbotin, D., Shprits, Y., & Ni, B. (2010). Three-dimensional VERB radiation belt simulations including mixed diffusion. *Journal of Geophysical Research*, 115, A03205. <https://doi.org/10.1029/2009JA015070>
- Takahashi, K., Lee, D., Merkin, V. G., Lyon, J. G., & Hartinger, M. D. (2016). On the origin of the dawn-dusk asymmetry of toroidal Pc5 waves. *Journal of Geophysical Research: Space Physics*, 121, 9632–9650. <https://doi.org/10.1002/2016JA023009>
- Takahashi, K., Yumoto, K., Claudepierre, S. G., Sanchez, E. R., Troshichev, O. A., & Janzhura, A. S. (2012). Dependence of the amplitude of Pc5-band magnetic field variations on the solar wind and solar activity. *Journal of Geophysical Research*, 117, A04207. <https://doi.org/10.1029/2011JA017120>
- Waskom, M., Botvinnik, O., O’Kane, D., Hobson, P., Ostblom, J., Lukauskas, S., et al. (2018). mwaskom/seaborn: v0.9.0 (July 2018). <https://doi.org/10.5281/ZENODO.1313201>
- Weisheimer, A., Palmer, T. N., & Doblas-Reyes, F. J. (2011). Assessment of representations of model uncertainty in monthly and seasonal forecast ensembles. *Geophysical Research Letters*, 38, L16703. <https://doi.org/10.1029/2011GL048123>
- Wright, A. N., Elsden, T., & Takahashi, K. (2018). Modeling the dawn/dusk asymmetry of field line resonances. *Journal of Geophysical Research: Space Physics*, 123, 6443–6456. <https://doi.org/10.1029/2018JA025638>
- Wright, A. N., & Rickard, G. J. (1995). ULF pulsations driven by magnetopause motions: Azimuthal phase characteristics. *Journal of Geophysical Research*, 100(A12), 23,703–23,710. <https://doi.org/10.1029/95JA01765>
- Wygant, J. R., Torbert, R. B., & Mozer, F. S. (1983). Comparison of S3-3 polar cap potential drops with the interplanetary magnetic field and models of magnetopause reconnection. *Journal of Geophysical Research*, 88(A7), 5727. <https://doi.org/10.1029/JA088iA07p05727>
- Zastenker, G. N., Nozdrachev, M. N., Němeček, Z., Šafránková, J., Paularena, K. I., Richardson, J. D., et al. (2002). Multispacecraft measurements of plasma and magnetic field variations in the magnetosheath: Comparison with Spreiter models and motion of the structures. *Planetary and Space Science*, 50(5–6), 601–612. [https://doi.org/10.1016/S0032-0633\(02\)00039-9](https://doi.org/10.1016/S0032-0633(02)00039-9)