

JGR Space Physics

RESEARCH ARTICLE

10.1029/2019JA027684

Key Points:

- We present a new model to forecast the maximum value of dB/dt over 20-min intervals at specific locations
- The model enhances the output of the physics-based Geospace model with a machine learning technique
- The model provides a probabilistic forecast of exceeding a predefined threshold at a given location

Correspondence to:

E. Camporeale,
enrico.camporeale@noaa.gov

Citation:

Camporeale, E., Cash, M. D., Singer, H. J., Balch, C. C., Huang, Z., & Toth, G. (2020). A gray-box model for a probabilistic estimate of regional ground magnetic perturbations: Enhancing the NOAA operational Geospace model with machine learning. *Journal of Geophysical Research: Space Physics*, 125, e2019JA027684. <https://doi.org/10.1029/2019JA027684>

Received 27 NOV 2019

Accepted 30 SEP 2020

Accepted article online 17 OCT 2020

A Gray-Box Model for a Probabilistic Estimate of Regional Ground Magnetic Perturbations: Enhancing the NOAA Operational Geospace Model With Machine Learning

E. Camporeale^{1,2} , M. D. Cash³, H. J. Singer³ , C. C. Balch³ , Z. Huang⁴, and G. Toth⁴ 

¹CIRES, University of Colorado Boulder, Boulder, CO, USA, ²Center for Mathematics and Computer Science (CWI), Amsterdam, Netherlands, ³Space Weather Prediction Center, NOAA, Boulder, CO, USA, ⁴Department of Climate and Space Sciences and Engineering, University of Michigan, Ann Arbor, MI, USA

Abstract We present a novel algorithm that predicts the probability that the time derivative of the horizontal component of the ground magnetic field dB/dt exceeds a specified threshold at a given location. This quantity provides important information that is physically relevant to geomagnetically induced currents (GICs), which are electric currents associated with sudden changes in the Earth's magnetic field due to space weather events. The model follows a “gray-box” approach by combining the output of a physics-based model with machine learning. Specifically, we combine the University of Michigan's Geospace model that is operational at the National Oceanic and Atmospheric Administration (NOAA) Space Weather Prediction Center, with a boosted ensemble of classification trees. We discuss the problem of recalibrating the output of the decision tree to obtain reliable probabilities. The performance of the model is assessed by typical metrics for probabilistic forecasts: Probability of Detection and False Detection, True Skill Statistic, Heidke Skill Score, and Receiver Operating Characteristic curve. We show that the ML-enhanced algorithm consistently improves all the metrics considered.

1. Introduction

Geomagnetically induced currents (GICs) represent one of the most severe risks posed by space weather events on our infrastructure on the ground, such as high-voltage power transmission systems. GICs are caused by sudden variations of the Earth's magnetic field that, through Faraday's law, induce a variation of the electric field (Boteler et al., 1998; Lanzerotti, 2001; Pirjola, 2007; Pirjola et al., 2000; Pulkkinen et al., 2005; Schrijver & Mitchell, 2013). The induced electric fields responsible for GICs can be estimated from the amplitude of the time derivative of magnetic fluctuations, often denoted as dB/dt , when combined with information of local Earth conductivity characteristics (Boteler & Pirjola, 1998; Horton et al., 2012; Ngwira et al., 2008; Pirjola, 2002; Viljanen et al., 2004). Hence, much attention has been dedicated to understanding and forecasting dB/dt (Viljanen, 1997, 2001).

Previous works on forecasting dB/dt can generally be divided into empirical and physics-based models. Empirical models exploit the statistical relationships between input quantities, such as solar wind observations recorded by satellites orbiting around L1 (first Lagrangian point) and the observed dB/dt at a specific station, with a typical time lag ranging between 15 and 60 min. Those statistical relationship can then be encoded into a regression model, in the form of a neural network, or a linear filter. Empirical models include Gleisner and Lundstedt (2001), Weigel et al. (2002, 2003), Wintoft (2005), Wintoft et al. (2005), Weimer (2013), Wintoft et al. (2015), Lotz and Cilliers (2015), among others.

On the other hand, physics-based models follow the evolution in time and space of the plasma and the electromagnetic field surrounding Earth and derive the ground magnetic field perturbation from physical laws. Typically, the spatial domain is divided in subregions, where the MHD approximation is used in the outer magnetosphere, while the inner magnetosphere and the transition to ionosphere are modeled by including kinetic processes. Examples of physics-based models that can, in principle, forecast dB/dt given the conditions of the solar wind observed at L1 are OpenGGCM (Open General Geospace Circulation Model; Raeder et al., 1998), GAMERA (Grid Agnostic MHD for Extended Research Applications; Zhang et al., 2019), LFM (Lyon-Fedder-Mobarry; Lanzerotti, 2004), and SWMF (Space Weather Modeling Framework; Tóth et al., 2005). Several works have assessed the ability of physics-based models to forecast

geomagnetic perturbations and more generally to recover plasma and field conditions as observed in the data (see, e.g., Gordeev et al., 2015; Jordanova et al., 2018; Pulkkinen et al., 2011; Rastätter et al., 2011, 2013; Welling, 2019; Welling & Ridley, 2010; Yu & Ridley, 2008). The validation and comparisons of different models for predicting dB/dt was specifically tackled in Pulkkinen et al. (2013) in order to support selecting a model to transition to operations at National Oceanic and Atmospheric Administration (NOAA)'s Space Weather Prediction Center (SWPC). As a result of that study, the University of Michigan's SWMF model, henceforth referred to as the Geospace model, was selected for transition to real-time operations.

In this paper we present a new model for predicting whether dB/dt will exceed given thresholds in a given time interval at specific locations. The model builds on the physics-based Geospace model. We show that the skill of the physics-based model can be considerably enhanced with a machine learning technique, improving all the performance metrics considered.

1.1. The Geospace Model at NOAA/SWPC

The Geospace model that runs operationally at NOAA/SWPC is a version of the Space Weather Modeling Framework developed by the University of Michigan (Tóth et al., 2005, 2012) that couples the following three physics domains. The outer magnetosphere is solved by BATS-R-US (Block-Adaptive-Tree Solar wind Roe-type Upwind Scheme) (Gombosi et al., 2004), the inner magnetosphere by the Rice Convection Model (RCM) (Toffoletto et al., 2003), and the ionosphere electrodynamics by the Ridley Ionosphere Model (RIM) (Ridley et al., 2004). A detailed description of the Geospace model and its modules can be found in (Pulkkinen et al., 2013; Tóth et al., 2014)

1.2. Prediction of dB/dt

In defining the problem, we follow the strategy introduced in Pulkkinen et al. (2013) and later adopted in Tóth et al. (2014) and Welling et al. (2017). Specifically, we define

$$dB/dt = \max_{\{t, t + \Delta t\}} \sqrt{(dB_n/dt)^2 + (dB_e/dt)^2} \quad (1)$$

as the maximum value of the time derivative of the horizontal magnetic field, over an interval Δt , where n and e denote the north and east components of the magnetic field, respectively. More specifically, we restrict the time interval to $\Delta t = 20$ min, and we cast the problem as a classification task. Namely, our model predicts the probability that dB/dt will exceed a given threshold at a given location, in a 20-min interval. We use overlapping intervals with a 1-min stride between subsequent intervals. Henceforth, we simply refer to dB/dt as defined in Equation 1.

As a proof-of-concept, we will show results for the following three magnetic stations: Fresno, California (Geomagnetic latitude: 43.12°N, operated by U.S. Geological Survey, USGS, code: FRN), Ottawa, Canada (Geomagnetic latitude: 54.88°N, operated by Geological Survey of Canada, GSC, code: OTT), Iqaluit, Canada (Geomagnetic latitude: 73.25°N, operated by GSC, code: IQA), hence testing our new method for low, middle, and high magnetic latitudes, respectively. The reported magnetic coordinates are derived from the International Geomagnetic Reference Field (IGRF) 12th generation (Thébault et al., 2015). The extension of this method to any other station is straightforward.

The need of enhancing a physics-based approach with machine learning can be appreciated by analyzing the accuracy of the Geospace model in predicting dB/dt . Figure 1 shows the number of instances of a given dB/dt value observed in the simulation (vertical axis) versus the corresponding value observed in the data (horizontal axis), both in logarithmic scale (FRN, OTT, and IQA stations shown in the left, middle, and right panels, respectively). Each column (i.e., a fixed observed value) is normalized to its maximum value. The statistics are computed over a 2-year interval (see below). The solid red line represents a perfect match between predicted and observed values. Figure 1 shows that the simulations tend to underestimate dB/dt for large values (particularly at high latitude) and overestimate it for small values (particularly at low latitude). Also, the range of observed and predicted values is dependent on the geomagnetic latitude, as expected.

The paper is divided as follows. Section 2 introduces the data used for this study and the corresponding time periods covered. Section 3 describes the methodology, including the machine learning technique, the performance metrics, and the features chosen in the model. Section 4 presents the results of the new model,




Figure 1. Two-dimensional histogram of the counts of dB/dt as obtained from the Geospace simulation (vertical axis) versus the corresponding measured values (horizontal axis). Both axes are in logarithmic scale, and the heat map is normalized columnwise with respect to the maximum value for each column, for better visualization. FRN, OTT, and IQA stations are respectively shown in the left, middle, and right panels.

comparing its performance with the output of the Geospace model alone, and emphasizes the probabilistic nature of the forecast. Finally, in section 5 we draw conclusions and make final remarks about future directions.

2. Data

The magnetic field historical records have been obtained by the International Real-time Magnetic Observatory Network (INTERMAGNET). The 1-min data in IAGA-2002 format were retrieved for the period 1 January 2001 to 5 May 2019 (<https://intermagnet.github.io/>) for the three stations (FRN, OTT, and IQA), consisting of about 9.45 M valid entries per station. The output of the Geospace model used for this work covers the time period 28 May 2017 to 5 May 2019, about 1,000,000 one-minute output values. In addition, we also use the output of the Geospace model discussed in Pulkkinen et al. (2013), evaluated over a time period covering six geomagnetic storms. Those simulation outputs are made available from the NASA Community Coordinated Modeling Center (CCMC) (https://ccmc.gsfc.nasa.gov/RoR_WWW/pub/dBdt/out/).

The Geospace model outputs the magnetic field at the location of the three stations at 1-min resolution. However, the inner boundary of the global MHD model is at $2.5 R_E$. Therefore, the magnetic perturbations at the geomagnetic observatories are calculated from the currents using Biot-Savart integrals, taking into account the following three contributions: the currents inside the BATS-R-US domain, the field-aligned currents in the gap region between 1 and $2.5 R_E$ radial distance, and the Pedersen and Hall currents in the ionosphere electrodynamics model RIM (Yu & Ridley, 2008).

In order to assess the accuracy of a trained machine learning model, it is important that the performance metrics are calculated on a portion of a data that has not been used for training (so-called unseen data). Moreover, when dealing with temporal data set, it is equally important that the training and test sets are temporally disjoint so to minimize the temporal correlations between the two and to ensure that the machine learning algorithm does actually learn some patterns and does not merely memorizes the training data. For all our experiments and results, we use as training set the period covering 28 May 2017 to 31 December 2018 and as test set the period 1 January 2019 to 5 May 2019. In addition, three of the six storm events used in Pulkkinen et al. (2013) have been added to the training set (events numbered 1, 3, and 4), and the rest have been used for testing. We have verified that the distribution of dB/dt is approximately equal between training and test sets. After excluding gaps in the data, the split results in about 450,000 data points in the training set and 90,000 in the test set.

3. Methodology

As mentioned in section 1, the goal of this work is not to predict the precise value of dB/dt for any given 20-min interval but rather to estimate the probability that a predefined threshold will be exceeded. Hence, the first task is to define such thresholds. In this paper, we slightly deviate from Pulkkinen et al. (2013), which focused on the following four thresholds: (0.3, 0.7, 1.1, 1.5) nT/s, independent of the station considered. Instead, we define thresholds specific for each location, by analyzing the overall distribution of

Table 1
Thresholds Considered for Each Station (in nT/s)

Station	60%	70%	80%	90%	95%
FRN	0.012	0.014	0.018	0.027	0.038
OTT	0.03	0.038	0.05	0.076	0.109
IQA	0.238	0.318	0.448	0.734	1.113

dB/dt observed in the INTERMAGNET data (~19 years of data) and choosing the following percentiles as thresholds: 60%, 70%, 80%, 90%, and 95%. The resulting thresholds are summarized in Table 1.

3.1. Metrics

The task under consideration is a probabilistic classification: For a given station the model outputs the probability that dB/dt will exceed a specified threshold value. Such a probabilistic outcome can be interpreted as a deterministic binary prediction (i.e., positive/negative) by simply assigning “positive” to all predictions above a certain probability, and “negative” otherwise. Once the probabilistic outcome is interpreted as a binary prediction, one can calculate the following quantities, defined over a certain number of predictions:

- P = total number of observed positives (event occurrences);
- N = total number of observed negatives (event non-occurrences);
- TP = True Positives: number of predicted positives that are observed positives;
- FP = False Positives: number of predicted positives that are observed negatives;
- TN = True Negatives: number of predicted negatives that are observed negatives;
- FN = False Negatives: number of predicted negatives that are observed positives;

and the following performance metrics:

- $TPR = TP/P$ = True Positive Rate (also called Probability of Detection, Sensitivity, Hit Rate);
- $FPR = FP/N$ = False Positive Rate (also called Probability of False Detection, False Alarm Rate);
- $TSS = TPR - FPR$ = True Skill Statistics.
- $HSS = 2(TP \cdot TN - FN \cdot FP) / (P(FN + TN) + N(TP + FP))$ = Heidke Skill Score.

The TPR measures the ability to find all positive events and a perfect classifier results in $TPR = 1$; the FPR measures the probability of wrongly classifying a negative as a positive, and a perfect classifier results in $FPR = 0$. Hence, TSS is a useful metric that combines both types of information and should be as close as possible to 1. Moreover, in a Receiver Operating Characteristic (ROC) curve, TPR and FPR are respectively on the vertical and horizontal axis, and TSS measures the distance to the diagonal (no-skill) line (Krzanowski & Hand, 2009). Finally, the HSS measures the skill of a method compared to a baseline represented by random chance. HSS has been used in Pulkkinen et al. (2013) and is used here for comparison with previous studies.

The baseline accuracy is represented by the True Skill Statistic (TSS) and the Heidke Skill Score (HSS) yielded by the Geospace model alone, that is, by calculating dB/dt directly from the simulation output and comparing to observations. Figure 2 shows the TSS (left) and HSS (right), where blue, red, and yellow lines are for FRN , OTT , and IQA stations, respectively. The scores are computed over all the data for which we have Geospace simulation outputs and are shown for different thresholds, represented on the horizontal axis in terms of their percentile calculated over ~19 years of observational data. One can notice that both scores are latitude dependent. Interestingly, while the TSS increases with higher percentiles (less frequent events) for FRN and OTT , the opposite is true for IQA . A different behavior is also noticeable with regards to the HSS score: FRN and OTT stations peak around the 70th percentile, while IQA peaks approximately at the 35th percentile. In general, Figure 2 shows that the Geospace model performs better at predicting large thresholds of dB/dt at middle and low latitudes than at high latitudes.

3.2. Machine Learning Classifier

A variety of methods exist in the Machine Learning arena to perform a probabilistic classification task. For this work, we have opted to use a boosted ensemble of classification trees. The method of choice is called RobustBoost (Freund, 2009). In this section we provide a short introduction and appropriate references.

Let us assume we want to assign a label $y \in \{0, 1\}$ to a data point $\mathbf{x} = \{x_1, x_2, \dots, x_D\}$, where D is the dimensionality of \mathbf{x} . The task is a supervised binary classification, meaning that we make use of a large data set of labeled examples to infer a pattern between the inputs \mathbf{x} and the binary outputs y that can be used to infer the label of new data points that have not been used to train the model. A decision tree is a simple method that recursively partitions the D -dimensional space of input variables one dimension at the time, thus creating a tree-like structure. In other words, by taking as decision boundaries hyperplanes defined by simple inequalities such as $x_i < c$, a decision tree divides the input space into a number of hypercubes where a given label is

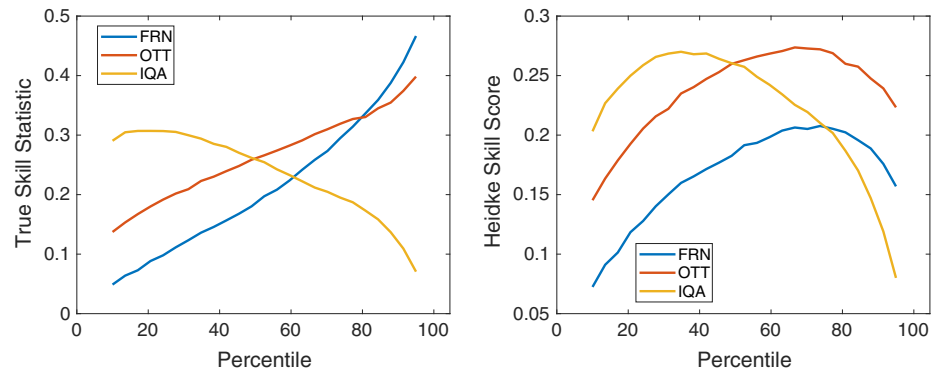


Figure 2. True Skill Statistic (left) and Heidke Skill Score (right) obtained from the predictions of the Geospace model, for different stations (in blue for FRN, red for OTT, and yellow for IQA), as functions of the different thresholds percentiles. The percentile are calculated on the distribution of observed dB/dt for a given station over a period of ~19 years of data.

assigned to all the data belonging to the same hypercube. Decision trees have the great advantage of being very transparent and easily interpretable. In fact, one can simply follow the tree structure from top to bottom to understand how a label is associated to a given data point. In order to choose where to set up a decision boundary (i.e. the value of the constant c) and along which variable, a partition criterion is followed. Two standard partition criteria are the Gini Index, and the Information Gain. The Gini Index measures the reduction in class impurity, which is defined as the probability that two randomly chosen data that belong to the same partition have different labels. At a given iteration when growing a tree, the best partition is the one that reduces such impurity, or in other words that minimizes the probability of a data point being mislabeled. The Information Gain is based on the entropy measured at each node and the optimal split is the one that minimizes the global entropy (or maximizes information). A reference monograph on decision trees is Breiman (1984). The accuracy of a classification tree can be improved by using a boosting strategy. Boosting refers to a class of algorithms that makes use of an ensemble of (not very accurate) predictions to produce a much more accurate one. The members of the ensemble are called weak learners and their weighted sum is referred to as strong learner. In the context of classification trees, the weak learners are represented by trees that are grown to only a few layers. One of the most successful and widely applied boosting techniques is *Adaboost* (short for Adaptive Boosting), introduced in the seminal paper by Freund and Schapire (1997). This is an algorithm that iteratively adds members to an existing ensemble. The newest member increasingly focuses on the data points that were misclassified by the previous members, and the weights of each member are iteratively adjusted. AdaBoost is typically less prone to overfitting than other algorithms, but it is very sensitive to outliers, because it will keep focusing on the few data points that are misclassified, eventually at the expense of the remaining data set. A modification of Adaboost that adds robustness to the algorithm (in the sense of not being so sensitive to outliers) is *RobustBoost*, being introduced in Freund (2009). RobustBoost can be intuitively understood as “giving up” on data points that are so far on the incorrect side of a decision boundary that they are unlikely to be correctly classified even after many iterations.

In this work we have used the MATLAB (R2019a) implementation of RobustBoost which is included in the Statistics and Machine Learning Toolbox. We have tested and compared the following boosting techniques: AdaBoost, GentleBoost, LogistBoost, RobustBoost, and Bagging, and although their results were comparable, RobustBoost is the algorithm that consistently yielded better results.

3.3. Feature Selection

In the machine learning jargon a feature is an explanatory variable that is used as an input for a given algorithm. The present work builds up on the idea presented in Tóth et al. (2014). The main finding was that a strong correlation exists between observed dB/dt and the observed maximum variation in the amplitude of the magnetic field, within the same 20-min interval. In fact, the correlation is almost linear when both quantities are expressed in their logarithm. Tóth et al. (2014) argued that the magnetic perturbations relative to




Figure 3. Two-dimensional histogram of the counts of the target variable dB/dt at the IQA station (vertical axis) and the six features described in section 3.3. Each heat map is normalized columnwise with respect to its maximum value.

the background dipole value obtained by the Geospace model simulations are much more reliable than the values of dB/dt computed directly from the field.

Feature selection refers to the procedure of selecting the most informative inputs used in a machine learning algorithm, making sure that the number of selected features is large enough for the algorithm to be accurate, but not too large, in order to prevent overfitting and for optimizing computational efficiency. The initial selection of features is done by a visual exploratory analysis of the correlation between candidate features and the target dB/dt . We have found that several quantities correlate well when plotted in logarithmic scale. As an example, we show in Figure 3 such correlations for the IQA station. The complete list of initially identified features is in Table 2. Here, the lead-time of the model forecasts ΔT are defined as the propagation time of the solar wind between the L1 point (where the solar wind is measured) and the outer boundary of the computational domain, approximately at the Earth's bow shock. In the OMNI data set used for this study solar wind quantities are conveniently time shifted to account for the propagation time between L1 and the bow shock. Hence, we have shifted back in time the measurements of Sym-H only, using the time shift provided by the OMNI data set (at 1-min resolution).

The procedure chosen to reduce the number of features is known as *backward elimination*, and it works as follows. First, we train a linear regression model using all the features listed in Table 2, thus assuming the output ($\log_{10}(dB/dt)$, measured from magnetometer data) to be a linear combination of the inputs, each weighted by a coefficient. The model returns both the values of the coefficients and their standard deviation. The t statistic (t value) is defined as the ratio between coefficients and their standard deviation. Inputs with low t value (in absolute value) are deemed noninformative. Therefore, we iteratively reduce the number of features by eliminating the one with the smallest t absolute value and we retrain a new linear model at each iteration with the remaining features. In this way we rank all the features listed in Table 2 (first column).

Table 2
Ranking of Features

Rank	Feature	Meaning	Data source	Time	Selected
1	$\log_{10}(dB/dt)$	target at previous time	magnetometer	$T - \Delta T$	yes
2	$\log_{10}(\max(Bn) - \min(Bn))$	Range of North component of magnetic field predicted by simulation	Geospace	T	yes
3	$\log_{10}(\max(Be) - \min(Be))$	Range of East component of magnetic field predicted by simulation	Geospace	T	yes
4	$\log_{10}(Bz)$	z component (GSM) of interplanetary magnetic field	OMNI data set	$T - \Delta T$	yes
5	$SymH$	Geomagnetic index Sym-H	OMNI data set	$T - \Delta T$	yes
6	dB/dt_{geo}	target predicted by simulation output	Geospace	T	no
7	$\log_{10}(\max(B) - \min(B))_{geo}$	Range of magnetic field amplitude predicted by simulation	Geospace	T	no
8	$\log_{10}(n)$	solar wind density	OMNI data set	$T - \Delta T$	no
9	$\log_{10}(\max(B) - \min(B))$	Range of magnetic field amplitude observed	magnetometer	$T - \Delta T$	no
10	$\log_{10}(E)$	Electric field	OMNI data set	$T - \Delta T$	no
11	$\log_{10}(V_x)$	x component of solar wind speed	OMNI data set	$T - \Delta T$	no

Note. T denotes the time at which dB/dt is predicted and ΔT is the solar wind propagation time.

Moreover, for each iteration we record the coefficient of determination R^2 as a metric for the goodness of fit. The final ranking of features is represented in Figure 4 that shows how the value of R^2 changes by increasingly adding features. The features on the horizontal axis are sorted in order of importance from left (most important) to right (less important) and each circle corresponds to a linear model that uses the named feature in addition to all the ones listed to its left. Not surprisingly, the past value of dB/dt is the most informative feature, yielding by itself a R^2 value of 0.805. However, the next two are features determined by the Geospace model output, namely, the difference between the maximum and minimum values of the North and East components of the magnetic field in a 20-min window. This justifies the gray-box philosophy of combining inputs from simulation outputs with past observations. On the basis of the backward elimination procedure, we decide to use the top five features of Figure 4, noticing that R^2 tends to plateau with more than five features. The procedure has been run on the combined training sets for all three stations. However, to avoid overfitting, at each iteration only 50% of the combined training set has been used to train the linear model. This explains the small fluctuations of R^2 during the plateau, that otherwise would be monotonically increasing, if subsequent models were trained on identical data.




Figure 4. Coefficient of determination R^2 for the linear models trained successively on a larger number of features. Each symbol represents a model trained with the feature reported on horizontal axis in addition to all the features to its left (see Table 2). The most important features are to the left and the less important to the right.

4. Results

In this section we show the results of our model in terms of the True Positive Rate (TPR, or probability of detection), False Positive Rate (FPR, or probability of false detection), TSS, and the HSS discussed in section 3.1. Essentially, a different classifier is trained for each station and each threshold. The proposed gray-box approach is compared against two alternative approaches: a white-box approach where one simply uses the value of dB/dt predicted by the Geospace model and a black-box approach where similar machine learning classifiers are trained, with the only difference of not using the inputs coming from the Geospace model. In other words, among the top five features listed in Table 2, the black-box models use only the three that do not come from Geospace output. Figure 5 shows the TPR (left) and FPR (right) for the three stations and as functions of the different threshold levels (see Table 1). Figure 6 shows TSS (left) and HSS (right) with the same format. One can notice that both black- and gray-box models largely outperform the corresponding white-box models. Moreover, although the results are dependent on stations and thresholds, the gray-box model further improves the




Figure 5. Probability of detection (left) and Probability of false detection (right) versus different thresholds (horizontal axis). Blue, red, and yellow lines denote respectively: a black-box model trained without using Geospace output, a gray-box model that uses both past observations and Geospace output, and a white-box model that uses only Geospace outputs.

black-box model, especially for large thresholds, which are the cases of most interest for space weather. On the other hand, whenever the white-box model yields large values for the probability of detection (e.g., for FRN station), the probability of false detection is also large, resulting in low values for both TSS and HSS.

4.1. Recalibration

As anticipated in section 1, the goal of this work is not to provide a binary classification, but rather to estimate the probability of exceeding predefined thresholds. In principle, classification trees can output probabilities, which are simply calculated as the observed ratio between positives and negatives on a given leaf (the final node on a decision tree) calculated over the whole training set. A well-known problem with classification trees is that such probabilities are often miscalibrated (Niculescu-Mizil & Caruana, 2007). Calibration refers to the consistency between the predicted probability assigned to an event and the actual frequency observed for that event. For instance, in the binary classification setting, if we collect all the instances in which a model predicts a probability p for a “positive” outcome (in our case, exceeding a threshold), that model is well-calibrated if on average a positive is actually observed with frequency p (the frequency being calculated over all those instances). One way to visualize the relationship between predicted probabilities and observed frequency is through a reliability diagram (DeGroot & Fienberg, 1983). To construct such diagram for binary classification, one discretizes the predicted probabilities in bins. For each bin, the average predicted frequency (horizontal axis) is plotted against the true fraction of positive cases in that bin (vertical axis). A perfect calibration will result in a diagonal straight line. Figure 7 shows the reliability diagrams for FRN, OTT, and IQA, respectively, in the top, middle, and bottom row. Each panel refers to a different threshold (see Table 1), and the blue circles represent the calibration of the boosted ensemble

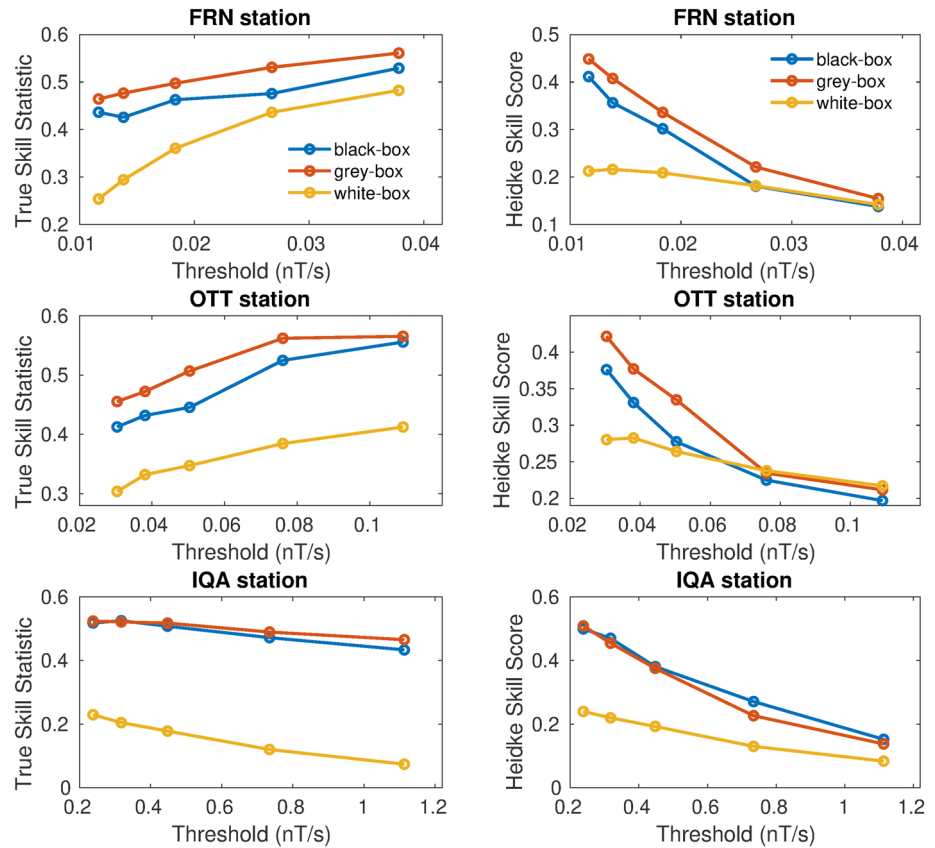


Figure 6. True Skill Statistic (left) and Heidke Skill Score (right) versus different threshold (horizontal axis). Blue, red, and yellow lines denote respectively: a black-box model trained without using Geospace output, a gray-box model that uses both past observations and Geospace output, and a white-box model that uses only Geospace outputs.

models, as trained by the MATLAB routine. One can clearly see that such predictions are miscalibrated. We apply a simple calibration strategy, where a mapping between old and new probabilities is derived by simply interpolating linearly the blue circles. For instance, a probability of 40% might be recalibrated to a new value of 30%. To perform recalibration fairly, we have derived the reliability diagram and the corresponding calibration map from the training set only. Figure 7 shows the reliability diagram calculated over the test set. The red diamonds represent the recalibrated reliability diagrams, that clearly suggest that all the models have been properly recalibrated. The recalibrated values are the ones that should be used to provide a probabilistic prediction.

4.2. ROC Curve

Another important diagnostic for a probabilistic model is the ROC curve. In order to interpret a probabilistic prediction in terms of true/false positives/negatives (see section 3.1), a probability threshold needs to be used to separate the predicted positives from the negatives. In the limit that such threshold is pushed to 0%, all the predictions become positives, which means that both the TPR and the FPR are equal to 1 (all positives are correctly predicted, but all negative are misclassified). In the opposite limit, when the threshold is 100% and all predictions are negative both TPR and FPR become equal to 0 (no positives are predicted, but all negatives are correctly predicted). The ROC curve is a continuous curve in the (FPR, TPR) space that connects these extreme scenarios (TPR = FPR = 1 and TPR = FPR = 0) by gradually changing the threshold from 0% to 100%. The optimal prediction is TPR = 1 and FPR = 0, and the optimal threshold is the point on the ROC curve with maximum distance from the diagonal. ROC curves for FRN, OTT, and IQA stations are shown in Figure 8, respectively, in the left, middle, and right panels. Different colors denote the five different thresholds, and an open circle represents the optimal values. Note that the TSS is the vertical distance

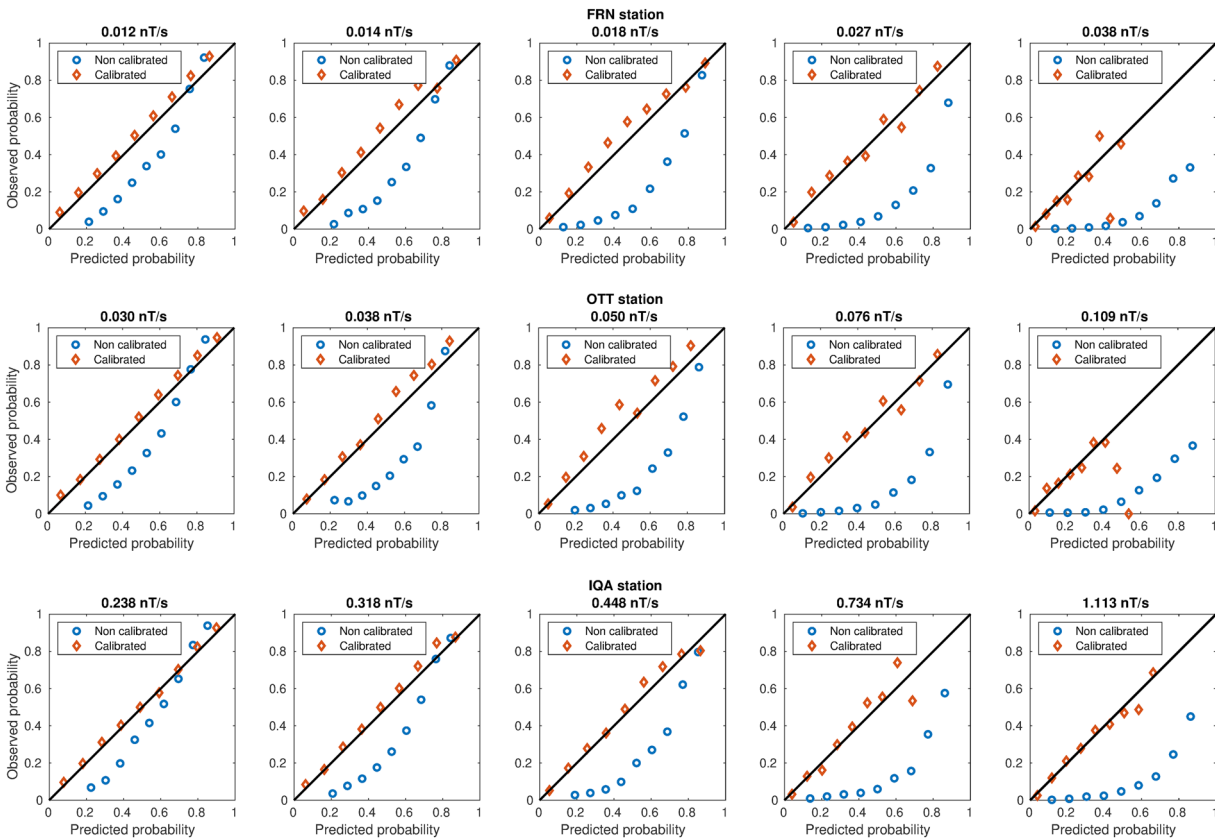


Figure 7. Reliability diagrams for different thresholds (increasing from left to right panels). Blue circles indicates the result of the noncalibrated models, and the red diamonds indicate the reliability achieved after recalibration. FRN, OTT, and IQA stations are shown in the top, middle, bottom rows, respectively.

between the ROC curve and the diagonal line ($TPR = FPR$), which represents no skill (i.e., a climatological forecast). The ROC curves demonstrate the general tendency of the models to improve their TSS for higher thresholds, as already shown in previous figures. Moreover, it is important to realize that the recalibration described in the previous section does not affect the ROC curve. In fact, by mapping old to new probabilities, the points on a given ROC curve get shifted along the same curve. In other words, what changes through recalibration is the value of the optimal threshold, but not the corresponding values of TPR, FPR, and Skill Scores. In practice, because the uncalibrated models tend to be overconfident (i.e., below the diagonal line in the reliability diagram), recalibration changes the optimal threshold from 50% to larger values. For instance, it can be that for a given model one needs to interpret as positives predictions with probabilities larger than 80% rather than 50%.




Figure 8. ROC curves (TPR vs. FPR) for different thresholds. Open circles indicate the optimal points along a given ROC curve. FRN, OTT, and IQA stations shown in the left, middle, and right panels, respectively.




Figure 9. Comparison of the True Skill Statistic (left) and Heidke Skill Score (right) for models using the output of the Geospace model alone (horizontal axis) versus the model presented in this paper (combining Geospace outputs with machine learning, vertical axis). The diagonal black line indicates no improvement.

5. Conclusions

We have developed a model that estimates the probability of dB/dt exceeding a given threshold, for three stations ranging from low, to mid and high latitudes (FRN, OTT, and IQA). Five different thresholds were chosen for each station, by calculating the 60th, 70th, 80th, 90th, and 95th percentiles on a long-span historic data set (~19 years). One of the crucial points of this work is that it combines a physics-based prediction provided by the Michigan Geospace model running at SWPC with a machine learning algorithm for binary classification, effectively following what is known as a gray-box approach (Camporeale, 2019; Camporeale et al., 2018). Indeed, we have shown that the Geospace model alone provides limited skills for predicting dB/dt , although we expect the model to improve over time by better capturing properties of the physical system. However, as already noted in Tóth et al. (2014), the maximum perturbation of the magnetic field within a 20-min interval correlates very strongly with dB/dt and hence it can be used as a predictor in a machine learning algorithm. The chosen machine learning algorithm is an ensemble of classification trees, adaptively boosted via RobustBoost (Freund, 2009), and the performance metrics that we have analyzed are the TPR (or probability of detection), FPR (or probability of false detection), TSS, and HSS. Finally, we have discussed the issue of re-calibration and the ROC curve relative to all models.

Overall the gray-box approach proposed in this paper consistently enhances the results of the corresponding white-box approach, where one would directly take the results of the Geospace model as predictors of dB/dt . Indeed, Figure 9 summarizes the findings of previous Figures by comparing the TSS(left panel) and the HSS (right panel) of the Geospace model alone (horizontal axis) against the corresponding results applying machine learning (vertical axis). Different symbols are for the three different stations, and the region above the diagonal black solid line denotes an improvement.

The new model will be a valuable addition to the operational capabilities of the NOAA's Space Weather Prediction Center. It will be straightforward to extend the model including several stations spanning a range of latitudes and longitudes. We are currently investigating what is the optimal strategy to represent in a compact graphical display the probabilistic outcomes for several stations and several thresholds, such that the SWPC forecasters can extract valuable real-time information on a regional scale and experiment how to incorporate such information in their forecast.

Data Availability Statement

The INTERMAGNET data used for this study are publicly available online (<https://intermagnet.github.io/>). The Geospace model outputs discussed in Pulkkinen et al. (2013) are made available by the NASA Community Coordinated Modeling Center (CCMC) (https://ccmc.gsfc.nasa.gov/RoR_WWW/pub/dBdt/out/). All the data and codes are available online (at <https://doi.org/10.5281/zenodo.4075025>).

Acknowledgments

The results presented in this paper rely on the data collected at Fresno (FRN), Ottawa (OTT), and Iqaluit (IQA) geomagnetic observatories. We thank the U.S. Geological Survey and Natural Resources Canada Geomagnetism Programs for supporting their operation and INTERMAGNET for promoting high standards of magnetic observatory practice (www.intermagnet.org). GT and ZH are partially funded by the NASA Heliophysics DRIVE Science Center (SOLSTICE) at the University of Michigan under grant NASA 80NSSC20K0600.

References

Boteler, D., & Pirjola, R. (1998). The complex-image method for calculating the magnetic and electric fields produced at the surface of the Earth by the auroral electrojet. *Geophysical Journal International*, 132(1), 31–40.

Boteler, D., Pirjola, R., & Nevanlinna, H. (1998). The effects of geomagnetic disturbances on electrical systems at the Earth's surface. *Advances in Space Research*, 22(1), 17–27.

Breiman, L., Friedman, J., Olshen, R. A., & Stone, C. J. (1984). *Classification and regression trees*. Boca Raton, FL: CRC Press.

Camporeale, E. (2019). The challenge of machine learning in space weather nowcasting and forecasting. *Space Weather*, 17, 1166–1207. <https://doi.org/10.1029/2018SW002061>

Camporeale, E., Wing, S., & Johnson, J. (2018). *Machine learning techniques for space weather*. Amsterdam, Netherlands: Elsevier.

DeGroot, M. H., & Fienberg, S. E. (1983). The comparison and evaluation of forecasters. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 32(1–2), 12–22.

Freund, Y. (2009). A more robust boosting algorithm. *arXiv preprint arXiv:0905.2138*.

Freund, Y., & Schapire, R. E. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1), 119–139.

Gleisner, H., & Lundstedt, H. (2001). A neural network-based local model for prediction of geomagnetic disturbances. *Journal of Geophysical Research*, 106(A5), 8425–8433.

Gombosi, T. I., Powell, K. G., De Zeeuw, D. L., Clauer, C. R., Hansen, K. C., Manchester, W. B., et al. (2004). Solution-adaptive magneto-hydrodynamics for space plasmas: Sun-to-Earth simulations. *Computing in Science & Engineering*, 6(2), 14.

Gordeev, E., Sergeev, V., Honkonen, I., Kuznetsova, M., Rastätter, L., Palmroth, M., et al. (2015). Assessing the performance of community-available global MHD models using key system parameters and empirical relationships. *Space Weather*, 13, 868–884. <https://doi.org/10.1002/2015SW001307>

Horton, R., Boteler, D., Overbye, T. J., Pirjola, R., & Dugan, R. C. (2012). A test case for the calculation of geomagnetically induced currents. *IEEE Transactions on Power Delivery*, 27(4), 2368–2373.

Jordanova, V. K., Delzanno, G. L., Henderson, M. G., Godinez, H. C., Jeffery, C., Lawrence, E. C., et al. (2018). Specification of the near-Earth space environment with shields. *Journal of Atmospheric and Solar-Terrestrial Physics*, 177, 148–159.

Krzanowski, W. J., & Hand, D. J. (2009). *ROC curves for continuous data, Monographs on Statistics and Applied Probability*. Boca Raton, FL: Chapman and Hall/CRC.

Lanzerotti, L. J. (2001). Space weather effects on technologies. *Space Weather*, 125, 11–22.

Lanzerotti, L. J. (2004). The Lyon-Fedder-Mobarry (LFM) global MHD magnetospheric simulation code. *Journal of Atmospheric and Solar-Terrestrial Physics*, 66, 1333.

Lotz, S., & Cilliers, P. (2015). A solar wind-based model of geomagnetic field fluctuations at a mid-latitude station. *Advances in Space Research*, 55(1), 220–230.

Ngwira, C. M., Pulkkinen, A., McKinnell, L.-A., & Cilliers, P. J. (2008). Improved modeling of geomagnetically induced currents in the south African power network. *Space Weather*, 6, S11004. <https://doi.org/10.1029/2008SW000408>

Niculescu-Mizil, A., & Caruana, R. (2007). Obtaining calibrated probabilities from boosting. In *Evaluation Methods for Machine Learning II - Technical Report WS-07-05, Papers from the AAAI Workshop* (p. 28). Menlo Park, CA: The AAAI Press.

Pirjola, R. (2002). Review on the calculation of surface electric and magnetic fields and of geomagnetically induced currents in ground-based technological systems. *Surveys in Geophysics*, 23(1), 71–90.

Pirjola, R. (2007). Space weather effects on power grids. In V. Bothmer & I. A. Daglis (Eds.), *Space weather- physics and effects* (pp. 269–288). Berlin, Heidelberg: Springer-Verlag.

Pirjola, R., Boteler, D., Viljanen, A., & Amm, O. (2000). Prediction of geomagnetically induced currents in power transmission systems. *Advances in Space Research*, 26(1), 5–14.

Pulkkinen, A., Kuznetsova, M., Ridley, A., Raeder, J., Vapirev, A., Weimer, D., et al. (2011). Geospace environment modeling 2008–2009 challenge: Ground magnetic field perturbations. *Space Weather*, 9, S02004. <https://doi.org/10.1029/2010SW000600>

Pulkkinen, A., Lindahl, S., Viljanen, A., & Pirjola, R. (2005). Geomagnetic storm of 29–31 October 2003: Geomagnetically induced currents and their relation to problems in the Swedish high-voltage power transmission system. *Space Weather*, 3, S08C03. <https://doi.org/10.1029/2004SW000123>

Pulkkinen, A., Rastätter, L., Kuznetsova, M., Singer, H., Balch, C., Weimer, D., et al. (2013). Community-wide validation of geospace model ground magnetic field perturbation predictions to support model transition to operations. *Space Weather*, 11, 369–385. <https://doi.org/10.1002/swe.20056>

Raeder, J., Berchem, J., & Ashour-Abdalla, M. (1998). The geospace environment modeling grand challenge: Results from a global geospace circulation model. *Journal of Geophysical Research*, 103(A7), 14,787–14,797.

Rastätter, L., Kuznetsova, M., Glocer, A., Welling, D., Meng, X., Raeder, J., et al. (2013). Geospace environment modeling 2008–2009 challenge: Dst index. *Space Weather*, 11, 187–205. <https://doi.org/10.1002/swe.20036>

Rastätter, L., Kuznetsova, M., Vapirev, A., Ridley, A., Wiltberger, M., Pulkkinen, A., et al. (2011). Geospace environment modeling 2008–2009 challenge: Geosynchronous magnetic field. *Space Weather*, 9, S04005. <https://doi.org/10.1029/2010SW000617>

Ridley, A., Gombosi, T., & DeZeeuw, D. (2004). Ionospheric control of the magnetosphere: Conductance. *Annales Geophysicae*, 22, 567–584.

Schrijver, C. J., & Mitchell, S. D. (2013). Disturbances in the us electric grid associated with geomagnetic activity. *Journal of Space Weather and Space Climate*, 3, A19.

Thébault, E., Finlay, C. C., Beggan, C. D., Alken, P., Aubert, J., Barrois, O., et al. (2015). International geomagnetic reference field: The 12th generation. *Earth, Planets and Space*, 67(1), 79.

Toffoletto, F., Sazykin, S., Spiro, R., & Wolf, R. (2003). Inner magnetospheric modeling with the rice convection model. *Space Science Reviews*, 107(1–2), 175–196.

Tóth, G., Meng, X., Gombosi, T. I., & Rastätter, L. (2014). Predicting the time derivative of local magnetic perturbations. *Journal of Geophysical Research: Space Physics*, 119, 310–321. <https://doi.org/10.1002/2013JA019456>

Tóth, G., Sokolov, I. V., Gombosi, T. I., Chesney, D. R., Clauer, C. R., De Zeeuw, D. L., et al. (2005). Space weather modeling framework: A new tool for the space science community. *Journal of Geophysical Research*, 110, A12226. <https://doi.org/10.1029/2005JA011126>

Tóth, G., Van der Holst, B., Sokolov, I. V., De Zeeuw, D. L., Gombosi, T. I., Fang, F., et al. (2012). Adaptive numerical algorithms in space weather modeling. *Journal of Computational Physics*, 231(3), 870–903.

- Viljanen, A. (1997). The relation between geomagnetic variations and their time derivatives and implications for estimation of induction risks. *Geophysical Research Letters*, *24*(6), 631–634.
- Viljanen, A., Nevanlinna, H., Pajunpää, K., & Pulkkinen, A. (2001). Time derivative of the horizontal geomagnetic field as an activity indicator. *Annales Geophysicae*, *19*, 1107–1118.
- Viljanen, A., Pulkkinen, A., Amm, O., Pirjola, R., & Korja, T. (2004). Fast computation of the geoelectric field using the method of elementary current systems and planar Earth models. *Annales Geophysicae*, *22*, 101–113.
- Weigel, R., Klimas, A., & Vassiliadis, D. (2003). Solar wind coupling to and predictability of ground magnetic fields and their time derivatives. *Journal of Geophysical Research*, *108*(A7), 1298. <https://doi.org/10.1029/2002JA009627>
- Weigel, R., Vassiliadis, D., & Klimas, A. (2002). Coupling of the solar wind to temporal fluctuations in ground magnetic fields. *Geophysical Research Letters*, *29*(19), 1915. <https://doi.org/10.1029/2002GL014740>
- Weimer, D. R. (2013). An empirical model of ground-level geomagnetic perturbations. *Space Weather*, *11*, 107–120. <https://doi.org/10.1002/swe.20030>
- Welling, D. (2019). Magnetohydrodynamic models of B and their use in GIC estimates. In J. L. Gannon, A. Swidinsky, Z. Xu (Eds.), *Geomagnetically Induced Currents from the Sun to the Power Grid, AGU Geophysical Monograph Series* (pp. 43–65). Hoboken, NJ: John Wiley & Sons.
- Welling, D., Anderson, B., Crowley, G., Pulkkinen, A., & Rastätter, L. (2017). Exploring predictive performance: A reanalysis of the geospace model transition challenge. *Space Weather*, *15*, 192–203. <https://doi.org/10.1002/2016SW001505>
- Welling, D., & Ridley, A. (2010). Validation of SWMF magnetic field and plasma. *Space Weather*, *8*, S03002. <https://doi.org/10.1029/2009SW000494>
- Wintoft, P. (2005). Study of the solar wind coupling to the time difference horizontal geomagnetic field. *Annales Geophysicae*, *23*, 1949–1957.
- Wintoft, P., Wik, M., Lundstedt, H., & Eliasson, L. (2005). Predictions of local ground geomagnetic field fluctuations during the 7–10 November 2004 events studied with solar wind driven models. *Annales Geophysicae*, *23*, 3095–3101.
- Wintoft, P., Wik, M., & Viljanen, A. (2015). Solar wind driven empirical forecast models of the time derivative of the ground magnetic field. *Journal of Space Weather and Space Climate*, *5*, A7.
- Yu, Y., & Ridley, A. J. (2008). Validation of the space weather modeling framework using ground-based magnetometers. *Space Weather*, *6*, S0500. <https://doi.org/10.1029/2007SW000345>
- Zhang, B., Sorathia, K. A., Lyon, J. G., Merkin, V. G., Garretson, J. S., & Wiltberger, M. (2019). GAMERA: A three-dimensional finite-volume MHD solver for non-orthogonal curvilinear geometries. *The Astrophysical Journal Supplement Series*, *244*(1), 20.