

Forecasting Geoeffective Events from Solar Wind Data and Evaluating the Most Predictive Features through Machine Learning Approaches

SABRINA GUASTAVINO,^{1,2} KATSIARYNA BAHAMAZAVA,³ EMMA PERRACCHIONE,³ FABIANA CAMATTARI,^{1,2}
GIANLUCA AUDONE,³ DANIELE TELLONI,² ROBERTO SUSINO,² GIANALFREDO NICOLINI,² SILVANO FINESCHI,²
MICHELE PIANA,^{1,2} AND ANNA MARIA MASSONE¹

¹*MIDA, Dipartimento di Matematica, Università di Genova, via Dodecaneso 35 16146 Genova, Italy*

²*Istituto Nazionale di Astrofisica (INAF), Osservatorio Astrofisico di Torino*

³*Dipartimento di Scienze Matematiche Giuseppe Luigi Lagrange, Politecnico di Torino, Torino, Italy*

(Received March 18, 2024; Revised March 18, 2024; Accepted March 18, 2024)

Submitted to ApJ

ABSTRACT

This study addresses the prediction of geomagnetic disturbances by exploiting machine learning techniques. Specifically, the Long-Short Term Memory recurrent neural network, which is particularly suited for application over long time series, is employed in the analysis of in-situ measurements of solar wind plasma and magnetic field acquired over more than one solar cycle, from 2005 to 2019, at the Lagrangian point L1. The problem is approached as a binary classification aiming to predict one hour in advance a decrease in the SYM-H geomagnetic activity index below the threshold of -50 nT, which is generally regarded as indicative of magnetospheric perturbations. The strong class imbalance issue is tackled by using an appropriate loss function tailored to optimize appropriate skill scores in the training phase of the neural network. Beside classical skill scores, value-weighted skill scores are then employed to evaluate predictions, suitable in the study of problems, such as the one faced here, characterized by strong temporal variability. For the first time, the content of magnetic helicity and energy carried by solar transients, associated with their detection and likelihood of geo-effectiveness, were considered as input features of the network architecture. Their predictive capabilities are demonstrated through a correlation-driven feature selection method to rank the most relevant characteristics involved in the neural network prediction model. The optimal performance of the adopted neural network in properly forecasting the onset of geomagnetic storms, which is a crucial point for giving real warnings in an operational setting, is finally showed.

Keywords: solar storms; solar wind; geoeffectiveness; space weather; machine learning; neural networks; feature ranking

1. INTRODUCTION

Space Weather (SW) science studies how solar-terrestrial interactions affect the geospace environment (Pulkkinen 2007); specifically, it involves predicting major solar disturbances that can pose significant risks to terrestrial facilities, with potential economic and security implications. The SW domain thus covers the various physical processes involved in the transfer of energy from solar wind and events to the Earth system, such as magnetic reconnection, the generation of ring currents in the terrestrial magnetosphere, and the interaction of solar charged particles with the Earth's atmosphere (Chapman & Ferraro 1931; Dungey 1961; Fairfield & Cahill 1966). The most energetic events on the solar surface are flares (Piana et al. 2022), intense electromagnetic emissions that can accelerate particles to relativistic velocities. They are often followed by massive eruptive events of chromospheric and coronal material, known as

Coronal Mass Ejections (CMEs; [Webb & Howard 2012](#)), during which solar plasma and magnetic field are expelled into interplanetary space. Flares and CMEs, along with High-Speed Streams (HSS) of particles, and Corotating Interaction Regions (CIRs) and Heliospheric Current Sheet (HCS) crossings, can initiate geomagnetic storms that, when particularly intense, can have severe consequences for human activities and ground- and space-based infrastructures, such as telecommunications systems, satellite orbits, or power grids and pipelines.

Geomagnetic disturbances are usually detected by measuring perturbations in magnetospheric electric currents induced by solar storms. Specifically, the DST (Ap) geomagnetic index, derived from a network of ground-based magnetometer stations near the equator (at subauroral latitudes), is an estimate of the variations in the horizontal component of the ring currents circling the Earth in the equatorial plane (of the electric currents aligned with the Earth’s magnetic field in the auroral ionosphere). DST and Ap thus provide an assessment of the severity of geomagnetic storms at low and high latitudes, respectively. The DST index (hourly acquired) or the equivalent SYM-H (acquired at one minute) are those most commonly deployed to reveal geomagnetic storms and measure their degree of geoeffectiveness. These are classified as moderate, intense or extreme depending on whether the DST index falls below thresholds of -50 , -100 and -250 nT ([Cander & Mihajlovic 1998](#)).

Recently, machine learning and deep learning techniques are increasingly being employed in space weather applications ([Camporeale et al. 2018](#)), especially in the prediction of solar flares ([Bobra & Couvidat 2015](#); [Barnes et al. 2016](#); [Campi et al. 2019](#); [Georgoulis et al. 2021](#); [Guastavino et al. 2022a, 2023a](#); [Georgoulis et al. 2024](#)), the onset of CMEs and their arrival time to Earth ([Guastavino et al. 2023b](#); [Singh et al. 2023](#)). Such studies are based on remote observations of the Sun and its atmosphere and, more specifically, in identifying CMEs and extracting their morphological/dynamic properties (such as angular width, velocity, and acceleration) from time sequences of coronagraphic white-light images, such as those from the Large Angle Spectroscopic Coronagraph (LASCO), as in [Pricopi et al. \(2022\)](#); [Vourlidas et al. \(2019\)](#).

Less attention has been devoted, however, to predicting the degree of geoeffectiveness of solar events impinging on the Earth. In [Telloni et al. \(2023\)](#) the problem was addressed using interplanetary plasma and magnetic field measurements (and, in particular, the intensity and the normal component to the ecliptic plane of the magnetic field vector, and the bulk speed, temperature, and density of the solar wind plasma) acquired in situ at the L1 Lagrangian point, 1.5 million km from Earth in the Sunward direction, by resorting to the use of different types of neural networks. The recurrent architectures were found to be the best in predicting geomagnetic events associated with SYM-H values below -50 nT, achieving 94% (70%) accuracy when the prediction was made 1 (8) hours in advance.

Similarly to [Telloni et al. \(2023\)](#), this paper aims to forecast the degree of severity of CME-induced geomagnetic storms, but explores and extends its predictive capabilities, using more refined and potentially better approaches. More specifically, as in the previous work, the problem was tackled as a binary classification with the goal of forecasting 1-hour ahead a decrease in the SYM-H geomagnetic index below -50 nT, using the Long-Short Term Memory (LSTM) recurrent network (which was shown to have the best performance in [Telloni et al. 2023](#)) applied to long time series of solar wind data. However, unlike [Telloni et al. \(2023\)](#), in which the strong data imbalance (SYM-H < -50 nT only 2% of the time) potentially affecting the proper neural network training was addressed by exploiting data augmentation, in the present study this issue is faced by using an appropriate loss function designed to automatically optimize a suitable skill score for evaluating predictions in the case of highly imbalanced datasets ([Marchetti et al. 2022](#)). The advantage of this approach is to avoid data handling. The prediction performances are also evaluated against value-weighted skill scores ([Guastavino et al. 2022b](#)) that are more suited for forecasting over time ([Hu et al. 2022](#); [Guastavino et al. 2022c](#)) and will be shown to better evaluate performance in predicting the onset of geomagnetic storms. In addition, we have computed an estimate of the uncertainty of prediction by performing several runs of the neural network in the case of several generations of training, validation, and test sets. More importantly, in the present analysis, besides all interplanetary parameters (hereafter features) directly acquired by the instruments aboard the spacecraft, i.e., plasma and magnetic field measurements of the solar wind, some derived quantities, such as magnetic helicity (which has been shown to be critical for properly detecting CMEs at L1, e.g., [Telloni et al. 2019](#)) and solar wind transported energy (which has been shown to play a crucial role in the magnetospheric response to solar drivers [Telloni et al. 2020](#)), were also considered as inputs to the neural network. Finally, a correlation-driven feature selection method is here used to rank the most relevant, i.e., predictive, features involved in the neural network prediction model.

The layout of the paper is as follows: § 2 presents the solar wind dataset used as input for the neural network model; § 3 describes the machine learning approach and correlation-driven feature ranking method; § 4 shows the forecasting performances and the most predictive features. Our conclusions are offered in § 5.

2. DATASET

It is well known that a huge amount of data is required to train, validate and test neural networks. The dataset used in this paper consists of 7888319 one-minute acquisitions related to multi-spacecraft interplanetary parameters acquired at L1 and geomagnetic activity indices measured on ground. The time period spans the years 2005 to 2019, thus encompassing the entire solar cycle 24, and the descending (ascending) phase of the previous (subsequent) solar cycle 23 (25). The Operating Missions as a Node on the Internet (OMNI; King & Papitashvili 2005) database was used, which is intended specifically to study the effects of the solar wind variations on the Earth’s magnetosphere. In fact, the interplanetary measurements are artificially time-shifted forward as if they had been acquired at the same location as the magnetospheric indices (in other words, the distance of 1.5 million km separating L1 from Earth has been nullified by adding to the acquisition time of the heliospheric quantities the time it takes the solar wind to reach Earth). This allows the temporal correlations between solar and magnetospheric indices to be explored directly. Among all available interplanetary measurements, the following parameters were taken into account in the study: the magnitude and components of the magnetic field vector \mathbf{B} , the components of the flow velocity vector \mathbf{V} , along with the bulk speed U , density ρ , and temperature T of the solar wind plasma. Some quantities derived from spacecraft measurements were also considered. These are the magnetic helicity H_m , the kinetic E_k and magnetic energy E_m , and the total energy $E = E_k + E_m$. For a comprehensive discussion on how these quantities are estimated in the solar wind, the Reader is referred to Telloni et al. (2019) and Telloni et al. (2020). Here it is only reminded that magnetic helicity is a measure of the degree of twisting of the magnetic field lines and is therefore a clear signature of the CME-embedded flux rope (which is a helical structure and therefore can be seen as a region of high magnetic helicity). It follows that H_m is beneficial for a proper detection of CMEs. On the other hand, it appears evident that the more energetic the CME is (either due to its high velocity and/or the intensity of the associated magnetic field), the more severe the induced geomagnetic event is expected, since more energy will be transferred to the geospace. As a result, E is a key parameter for correct prediction of geoeffectiveness of solar storms. As for the assessment of geomagnetic activity, the SYM-H index was instead employed.

As mentioned above, the goal of this study is to predict whether, based on a history of features over the past 24 hours, in the next hour the SYM-H geomagnetic index will drop below the -50 nT threshold, customarily referred to for potential severe magnetospheric disturbances. From a purely computational perspective (i.e., to save machine resources), the time series are resampled at a one-hour resolution. This means that the total number of samples is thus reduced to just over 130000. Since the study is approached as a binary prediction problem, samples corresponding to time periods when SYM-H ≤ -50 nT are labelled with 1 or 0, respectively. As already noted in Telloni et al. (2023) and reported above, the dataset is highly imbalanced: only 2.53% of total samples are labeled with 1. The next section will present a suitable approach associated with the employment of an appropriate loss function to address this problem and enable proper training and validation of the neural network based on the use of suitable metrics for performance evaluation.

3. PREDICTION AND FEATURE RANKING

3.1. Assessment of results

In order to compare the performances of machine learning methods for binary classification problems, as the prediction of geoeffective events, the following points should be accounted for. First, the classification results should be evaluated by considering skill scores that are suitable for imbalanced data classification. Indeed, the percentage of geoeffective events is really small, as already pointed out in Section 2. Therefore, a chosen score needs to be capable of representing the performance of the classifier concerning the *small* positive class. The classical skill scores are computed on the entries of the so-called confusion matrix, which is defined as

$$\text{CM} = \begin{pmatrix} \text{TN} & \text{FP} \\ \text{FN} & \text{TP} \end{pmatrix}, \quad (1)$$

where the four entries are true positives (TPs), i.e. the number of samples labeled with 1 and correctly predicted as positive; true negatives (TNs), i.e. the number of samples labeled with 0 and correctly predicted as negative; false positives (FPs), i.e. the number of samples labeled with 0 incorrectly predicted as positive; and false negatives (FNs), i.e. the number of samples labeled with 1 and incorrectly predicted as negative. Among all possible skill scores, the True Skill Statistic (TSS) (Hanssen & Kuipers 1965), is less sensible to the class-imbalance ratio (Bloomfield et al.

2012) than others, and therefore, it is particularly suitable for evaluation of imbalanced classification tasks. It is defined as the balance between the true positive and true negative rates (named also sensitivity and specificity, respectively), as follows

$$\text{TSS}(\text{CM}) = \frac{\text{TP}}{\text{TP} + \text{FN}} + \frac{\text{TN}}{\text{FP} + \text{TN}} - 1 = \text{SENS}(\text{CM}) + \text{SPEC}(\text{CM}) - 1, \quad (2)$$

its values have range in the interval $[-1, 1]$, and the performance is optimal when $\text{TSS} = 1$, while $\text{TSS} > 0$ means that the rates of positive and negative events are mixed up.

Second, the evaluation of binary predictions performed over time should take into account the forecast value, measured in terms of its usefulness to an operational setting to support the user while making a decision, as the importance in predicting the starting time of a geomagnetic storm. Value-weighted skill scores introduced by Guastavino et al. (2022b) that better account the intrinsic dynamical nature of forecasting problems are defined on a value-weighted confusion matrix that assigns different weights to FPs (denoted by wFPs) and FNs (denoted by wFNs) in such a way to account for the distribution of predictions along time with respect to the actual occurrences. By denoting the value-weighted confusion matrix as

$$\text{wCM} = \begin{pmatrix} \text{TN} & \text{wFP} \\ \text{wFN} & \text{TP} \end{pmatrix}, \quad (3)$$

predictions are assessed by computing the value-weighted TSS (wTSS) defined as follows

$$\text{wTSS} = \text{TSS}(\text{wCM}). \quad (4)$$

Finally, the splitting strategy between training, validation and test sets should take into account the rare-event nature of the problem, by maintaining uniformly the percentage of 1-labelled samples and the temporal distribution of events in order to not mix temporally samples in training, validation and test sets. Furthermore, the splitting strategy should be repeated several times in order to achieve some statistical significance. Therefore, many classification tests should be carried out by generating different triples of training, validation and test sets by maintaining the positive percentage of samples and not drawing samples completely randomly between training and test with respect to time.

3.2. Recurrent neural network

As neural network, we adopted the long short-term memory (LSTM) network (Hochreiter & Schmidhuber 1997), which is the most widely adopted type of recurrent neural network, able to process sequential data by solving the well-known short-term memory problem of basic recurrent architectures. In the experiments we set the LSTM units equal to 72 followed by a dense layer of 64 neurons where the rectified linear unit (ReLU) activation function is adopted and a final dense sigmoid unit drives the output of the network to be in the interval $[0, 1]$, in order to perform binary prediction. The LSTM network is trained for 100 epochs by taking batch size of 256 samples, the *Adam* optimizer (Kingma & Ba 2015) is adopted for the optimization process with learning rate equal to 10^{-4} .

In order to face the class-imbalance issue we adopted a suitable loss function which is designed to optimize an appropriate score. This strategy is introduced in Marchetti et al. (2022) and it is based on a probabilistic definition $\overline{\text{CM}}$ of the classical confusion matrix, depending on a chosen cumulative density function (cdf) on $[0, 1]$ for the threshold parameter τ which separates 0 and 1. Let $(x_i, y_i)_{i=1}^n$ be input-label samples where $y_i \in \{0, 1\}$ is the actual label associated to the sample x_i and let $f(x_i) \in (0, 1)$ is the probability outcome of the neural network f on the sample x_i . We defined the probabilistic confusion matrix $\overline{\text{CM}}$ as

$$\overline{\text{CM}} = \begin{pmatrix} \overline{\text{TN}} & \overline{\text{FP}} \\ \overline{\text{FN}} & \overline{\text{TP}} \end{pmatrix}, \quad (5)$$

characterized by the following entries

$$\begin{aligned} \overline{\text{TP}} &= \sum_{i=1}^n y_i f(x_i), & \overline{\text{TN}} &= \sum_{i=1}^n (1 - y_i)(1 - f(x_i)) \\ \overline{\text{FP}} &= \sum_{i=1}^n (1 - y_i) f(x_i), & \overline{\text{FN}} &= \sum_{i=1}^n y_i (1 - f(x_i)). \end{aligned} \quad (6)$$

We chose the Score-Oriented Loss function (SOL) based on the TSS, which is defined as follows

$$\ell_{\text{TSS}} := -\text{TSS}(\overline{\text{CM}}). \quad (7)$$

The main advantage consists in an automatic optimization of the desired skill score during the training phase without the need of a posteriori optimization of the threshold.

3.3. Correlation-driven feature ranking

Feature ranking methods are commonly employed to identify a reduced subset of highly predictive features and to assess the relevance of physical attributes. One common technique consists in evaluating how each feature impacts the predictions according to some permutation importance score. Once a model is fitted via a training set, the permutation importance algorithm is implemented to evaluating how the accuracy of the prediction changes when a single feature is randomly shuffled in the validation dataset (Breiman 2001). Indeed, when a feature is shuffled, its relevance increases as the accuracy score on the prediction decreases. Nevertheless, the permutation feature importance algorithm does not take into account the correlation between features. Therefore, we consider a correlation-driven feature importance method inspired by Kaneko (2022), which includes the absolute correlation coefficients between features in the permutation process. Hence, in presence of strongly correlated features, such a method leads to more stable and reliable feature ranking schemes. In the following we summarize the steps of the correlation-driven feature importance algorithm:

- Train a neural network \hat{f} .
- Compute the prediction of the trained neural network \hat{f} on the validation data, denoted as $\mathbf{Z}_{val} = \{(\mathbf{X}_{val}, \mathbf{y}_{val})\}$ and compute the reference score s_{val} , i.e. the score computed between the prediction $\hat{f}(\mathbf{X}_{val})$ and \mathbf{y}_{val} .
- Calculate the correlation coefficients between all the features by following the procedure in Kaneko (2022); in particular the correlation coefficient between two features is set to 0 if there is the possibility of chance correlation.
- For each feature j , i.e. the j -th column of validation data \mathbf{X}_{val} , and for each repetition $k \in \{1, \dots, K\}$, randomly shuffle the j -th feature of the validation dataset \mathbf{X}_{val} , and for each feature $l \neq j$ for which the correlation value $c_{j,l}$ is higher than 0 randomly sample the column l of \mathbf{X}_{val} (without duplication with a probability of $c_{j,l}$) to generate a corrupted version of the validation dataset, denoted as $\mathbf{X}_{val}^{j,k}$.
- Compute the score $s_{j,k}$ between the prediction on the corrupted data $\mathbf{X}_{val}^{j,k}$, i.e. $\hat{f}(\mathbf{X}_{val}^{j,k})$, and \mathbf{y}_{val} .
- Evaluate the importance \mathcal{I}_j of the j -th feature by computing the difference between the reference score and the mean score obtained on the corrupted validation datasets, i.e.:

$$\mathcal{I}_j = s_{val} - \frac{1}{K} \sum_{k=1}^K s_{j,k}. \quad (8)$$

Such a procedure allows obtaining a feature ranking: the higher is \mathcal{I}_j , the higher is the contribution of the feature in the prediction.

4. RESULTS

The LSTM network described in section 3.2 has been applied to time series of features described in Section 2. First a splitting strategy based on stratified K-fold splitting allows the generation of training, validation and test set that share the same rate of about 2.5% of 1-labelled samples, leading to a training and validation sets which the sum represents about the 75% of the total number of sample and a test set representing the remaining 25% of samples. We investigated how the prediction performances change if the SYM-H of the past 24-hour with cadence 1 hour is added as additional feature to the considered solar wind features, i.e. B, Bx, By, Bz (which represent the absolute value of the magnetic field vector \mathbf{B} and the three components of \mathbf{B}), V, Vx, Vy, Vz (which represent the absolute value of the velocity vector \mathbf{V} and the three components of \mathbf{V}), ρ (density), T (temperature), E_k (kinetic energy), E_m (magnetic energy), E (total energy) and H_m (magnetic helicity). When the network is trained we applied the correlation-driven feature importance method described in section 3.3, in order to rank features accordingly to their predictive capabilities and we selected the first 10 features which are associated to a meaningful importance value at least higher then 0.1. In table 1 we report the ranking of features with the associated importance value both in the case the SYM-H is included or not between features. In detail, we compute the confusion matrix CM and some common skill scores as the Heidke Skill Score (HSS) (Heidke 1926), the sensitivity (SENS), the specificity (SPEC) defined in equation (2), the F1 score (F1), which is the harmonic mean of precision and sensitivity, and balanced accuracy (BA), which is the arithmetic mean between sensitivity and specificity. We noticed that the most predictive solar wind features are almost the same,

both when SYM-H is included or not between features at least of permutation. In particular the energies, as the Total Energy is the third one in both rankings: this is coherent with the analysis conducted between the DST and energy in [Telloni et al. \(2020\)](#). Further, as we expected, features that are not correlated with the occurrence of geomagnetic storms have a negligible importance value less than 0.1. In [table 2](#) we report the prediction performances on the test set when all the features were used and if the SYM-H is included in the list of features compared with the performance when only the subset of predictive features is considered. We noticed that the performances are higher when the subset of the most predictive features is adopted: this confirms the importance of a feature selection procedure also in the prediction phase. Finally in order to assess statistical robustness we reported results when 5 different splitting of training, validation and test sets are considered. In [Table 3](#), we reported the TSS and value-weighted TSS (wTSS) among the 5 generated test sets. We pointed out that when SYM-H is not used as feature the performances are low: this is due to the fact that when a geomagnetic storm is in place then, the SYM-H is below -50nT and such information helps the network to learn that if the previous hour the SYM-H is below -50nT the probability that in the next hour is below this threshold is high. Therefore, in [figures 1 and 2](#) we report the distribution of predictions along time in correspondence of two different temporal windows which include some geomagnetic storms. When SYM-H is not used as additional feature we noticed that more false positives are present. However, from all panels it is clear that the starting time of the geomagnetic storm is well predicted even when SYM-H is not used as a feature: this confirms the importance of solar wind features that give information about the possible occurrence of a geomagnetic storm, whether caused by a coronal mass ejection or other events such as high speed streams.

Table 1. Feature rankings with respect to the correlation-driven feature importance method shown in [section 3.3](#) in the case SYM-H is included between features or not.

| Included SYM-H | | Not included SYM-H | |
|-----------------------------|------------------|-----------------------------|------------------|
| Feature | Importance value | Feature | Importance value |
| 1) SYM-H | 0.5577 | 1) B | 0.2936 |
| 2) B | 0.3392 | 2) Bz | 0.2605 |
| 3) Total Energy | 0.3334 | 3) Total Energy | 0.2175 |
| 4) Vx | 0.3256 | 4) Magnetic Helicity | 0.2042 |
| 5) V | 0.3162 | 5) Vx | 0.1902 |
| 6) Temperature | 0.287 | 6) Magnetic Energy | 0.1901 |
| 7) Magnetic Energy | 0.2793 | 7) Temperature | 0.18 |
| 8) Magnetic Helicity | 0.2628 | 8) V | 0.1783 |
| 9) Kinetic Energy | 0.2427 | 9) Kinetic Energy | 0.1604 |
| 10) Bz | 0.2303 | 10) Density | 0.1451 |
| 11) Density | 0.2044 | 11) Vy | < 0.1 |
| 12) Vy | < 0.1 | 12) By | < 0.1 |
| 13) By | < 0.1 | 13) Bx | < 0.1 |
| 14) Vz | < 0.1 | 14) Vz | < 0.1 |
| 15) Bx | < 0.1 | - | - |

5. CONCLUSIONS

In this work, the performance of the LSTM recurrent neural architecture in providing one-hour advance alerts of critical SYM-H values less than -50 nT and thus indicative of geomagnetic disturbances was estimated. It turns out that its predictive capability reaches up to 91% and 82% in TSS in the case of considering or not SYM-H itself as an input feature of the neural network. It appears evident that in the former case the neural network is able to predict the activity status of the Earth’s magnetosphere even after the onset of the geomagnetic storm, that is, to estimate its recovery phase toward a quiet condition. This results in higher performance, as it is easier to also forecast the duration of the recovery phase of a solar storm rather than just its onset, which is precisely what LSTM does when not having the value of SYM-H as input. From an operational point of view, it is also worth mentioning that the solar wind data used in the analysis were artificially time-shifted as if they had been acquired at the same location as the geomagnetic indices. However, if the tool worked in real time in forecasting the value of SYM-H in the hour after the


Table 2. Comparison of the performances on the test set between using all features or the selected features with the correlation-driven feature importance method.

| CM | Included SYM-H | | | | Not included SYM-H | | | |
|------|-----------------|--------------------|-------------------|------------------|--------------------|------------------|-------------------|--------------------|
| | All features | | Selected features | | All features | | Selected features | |
| | TP=629 FP=74 | FN=204 TN=31599 | TP=707 FP=430 | FN=124 TN=615 | TP=495 FP=36 | FN=336 TN=668 | TP=550 FP=430 | FN=281 TN=31599 |
| TSS | 0.742 | | 0.8412 | | 0.5766 | | 0.633 | |
| HSS | 0.655 | | 0.7593 | | 0.4963 | | 0.4622 | |
| SENS | 0.755 | | 0.8508 | | 0.5957 | | 0.6619 | |
| SPEC | 0.987 | | 0.9904 | | 0.9809 | | 0.9715 | |
| F1 | 0.665 | | 0.766 | | 0.511 | | 0.48 | |
| BA | 0.871 | | 0.9206 | | 0.7883 | | 0.8267 | |


Table 3. Performances on 5 test sets in case the SYM-H is included or not in the subset of features.

| Splitting | TSS | | wTSS | |
|----------------|--------------------|--------------------|--------------------|--------------------|
| | Included SYM-H | Not included SYM-H | Included SYM-H | Not included SYM-H |
| 1 | 0.9073 | 0.82029 | 0.8964 | 0.7597 |
| 2 | 0.8654 | 0.7146 | 0.8503 | 0.6166 |
| 3 | 0.9082 | 0.7791 | 0.8973 | 0.6743 |
| 4 | 0.8442 | 0.6296 | 0.7838 | 0.4844 |
| 5 | 0.846 | 0.6334 | 0.8164 | 0.4757 |
| Mean \pm std | 0.8742 ± 0.028 | 0.7154 ± 0.076 | 0.8164 ± 0.044 | 0.6022 ± 0.109 |

acquisition of solar measurements at L1, the time required for the solar plasma to reach Earth, i.e., 30 minutes (60 minutes) for a bulk speed of 800 (400) km s^{-1} , would also have to be accounted for. This means that the prediction of geomagnetic disturbances could be provided with 91% efficiency up to 2 hours before the onset of the geomagnetic storm, ranking the present approach among those currently available that provide at the same time earlier and more accurate warnings.




(a) SYM-H not included in features.




(b) SYM-H included in features.

Figure 1. Predictions over time on a temporal window of the test set of splitting 1: the top panel represents the prediction when SYM-H is not included in the subset of features, whereas bottom panel represents the one when SYM-H is included between features.



(a) SYM-H not included in features



(b) SYM-H included in features

Figure 2. Predictions over time on a temporal window of the test set of splitting 1: the top panel represents the prediction when SYM-H is not included in the subset of features, whereas bottom panel represents the one when SYM-H is included between features.

ACKNOWLEDGEMENT

SG was supported by the Programma Operativo Nazionale (PON) “Ricerca e Innovazione” 2014–2020; AMM, EP, RS, DT, GN acknowledge the support of the Fondazione Compagnia di San Paolo within the framework of the Artificial Intelligence Call for Proposals, AIxtreme project (ID Rol: 71708). AMM is also grateful to the HORIZON Europe ARCAFF Project, Grant No. 101082164.

REFERENCES

- Barnes, G., Leka, K. D., Schrijver, C. J., et al. 2016, *ApJ*, 829, 89, doi: [10.3847/0004-637X/829/2/89](https://doi.org/10.3847/0004-637X/829/2/89)
- Bloomfield, D. S., Higgins, P. A., McAteer, R. J., & Gallagher, P. T. 2012, *The Astrophysical Journal Letters*, 747, L41
- Bobra, M. G., & Couvidat, S. 2015, *ApJ*, 798, 135, doi: [10.1088/0004-637X/798/2/135](https://doi.org/10.1088/0004-637X/798/2/135)
- Breiman, L. 2001, *Machine learning*, 45, 5
- Campi, C., Benvenuto, F., Massone, A. M., et al. 2019, *The Astrophysical Journal*, 883, 150
- Camporeale, E., Wing, S., Johnson, J., Jackman, C. M., & McGranaghan, R. 2018, *Space Weather*, 16, 2, doi: [10.1002/2017SW001775](https://doi.org/10.1002/2017SW001775)
- Cander, L. R., & Mihajlovic, S. J. 1998, *J. Geophys. Res.*, 103, 391, doi: [10.1029/97JA02418](https://doi.org/10.1029/97JA02418)
- Chapman, S., & Ferraro, V. C. A. 1931, *Terrestrial Magnetism and Atmospheric Electricity*, 36, 77, doi: [10.1029/TE036i002p00077](https://doi.org/10.1029/TE036i002p00077)
- Dungey, J. W. 1961, *PhRvL*, 6, 47, doi: [10.1103/PhysRevLett.6.47](https://doi.org/10.1103/PhysRevLett.6.47)
- Fairfield, D. H., & Cahill, L. J., Jr. 1966, *J. Geophys. Res.*, 71, 155, doi: [10.1029/JZ071i001p00155](https://doi.org/10.1029/JZ071i001p00155)
- Georgoulis, M. K., Bloomfield, D. S., Piana, M., et al. 2021, *Journal of Space Weather and Space Climate*, 11, 39, doi: [10.1051/swsc/2021023](https://doi.org/10.1051/swsc/2021023)
- Georgoulis, M. K., Yardley, S. L., Guerra, J. A., et al. 2024, *Advances in Space Research*
- Guastavino, S., Marchetti, F., Benvenuto, F., Campi, C., & Piana, M. 2022a, *Astronomy & Astrophysics*, 662, A105
- . 2023a, *Frontiers in Astronomy and Space Sciences*, 9, doi: [10.3389/fspas.2022.1039805](https://doi.org/10.3389/fspas.2022.1039805)
- Guastavino, S., Piana, M., & Benvenuto, F. 2022b, *IEEE Transactions on Neural Networks and Learning Systems*
- Guastavino, S., Piana, M., Tizzi, M., et al. 2022c, *Scientific Reports*, 12, 20049
- Guastavino, S., Candiani, V., Bemporad, A., et al. 2023b, *The Astrophysical Journal*, 954, 151, doi: [10.3847/1538-4357/ace62d](https://doi.org/10.3847/1538-4357/ace62d)
- Hanssen, A., & Kuipers, W. 1965, *On the Relationship Between the Frequency of Rain and Various Meteorological Parameters: (with Reference to the Problem of Objective Forecasting)*, Koninkl. Nederlands Meteorologisch Instituut. Mededelingen en Verhandelingen (Staatsdrukkerij- en Uitgeverijbedrijf). <https://books.google.it/books?id=nTZ8OgAACAAJ>
- Heidke, P. 1926, *Geogr. Ann.*, 8, 301
- Hochreiter, S., & Schmidhuber, J. 1997, *Neural Computation*, 9, 1735, doi: [10.1162/neco.1997.9.8.1735](https://doi.org/10.1162/neco.1997.9.8.1735)
- Hu, A., Shneider, C., Tiwari, A., & Camporeale, E. 2022, *Space Weather*, 20, e2022SW003064
- Kaneko, H. 2022, *Analytical Science Advances*, 3, 278
- King, J. H., & Papitashvili, N. E. 2005, *Journal of Geophysical Research (Space Physics)*, 110, A02104, doi: [10.1029/2004JA010649](https://doi.org/10.1029/2004JA010649)
- Kingma, D. P., & Ba, J. 2015, in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, ed. Y. Bengio & Y. LeCun. <http://arxiv.org/abs/1412.6980>
- Marchetti, F., Guastavino, S., Piana, M., & Campi, C. 2022, *Pattern Recognition*, 132, 108913
- Piana, M., Emslie, A. G., Massone, A. M., & Dennis, B. R. 2022, *Hard X-Ray Imaging of Solar Flares*, Vol. 164 (Springer)
- Pricopi, A.-C., Paraschiv, A. R., Besliu-Ionescu, D., & Marginean, A.-N. 2022, *The Astrophysical Journal*, 934, 176
- Pulkkinen, T. 2007, *Living Reviews in Solar Physics*, 4, 1, doi: [10.12942/lrsp-2007-1](https://doi.org/10.12942/lrsp-2007-1)
- Singh, T., Benson, B., Raza, S. A. Z., et al. 2023, *ApJ*, 948, 78, doi: [10.3847/1538-4357/acc10a](https://doi.org/10.3847/1538-4357/acc10a)
- Telloni, D., Antonucci, E., Bemporad, A., et al. 2019, *ApJ*, 885, 120, doi: [10.3847/1538-4357/ab48e9](https://doi.org/10.3847/1538-4357/ab48e9)
- Telloni, D., Carbone, F., Antonucci, E., et al. 2020, *ApJ*, 896, 149, doi: [10.3847/1538-4357/ab91b9](https://doi.org/10.3847/1538-4357/ab91b9)
- Telloni, D., Zhao, L., Zank, G. P., et al. 2020, *The Astrophysical Journal Letters*, 905, L12
- Telloni, D., Schiavo, M. L., Magli, E., et al. 2023, *ApJ*, 952, 111, doi: [10.3847/1538-4357/acdea](https://doi.org/10.3847/1538-4357/acdea)

Vourlidas, A., Patsourakos, S., & Savani, N. P. 2019, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 377, 20180096, doi: [10.1098/rsta.2018.0096](https://doi.org/10.1098/rsta.2018.0096)

Webb, D. F., & Howard, T. A. 2012, Living Reviews in Solar Physics, 9, 3, doi: [10.12942/lrsp-2012-3](https://doi.org/10.12942/lrsp-2012-3)